

VOLUME 13 PART 3 OCTOBER 1994

BRITISH TELECOMMUNICATIONS ENGINEERING



Included in this Issue

*Building on the Architectural
Framework*

*New Structure for London's PSTN
'Value' of the Telecommunications
Engineer*



**The Journal of The Institution of
British Telecommunications Engineers**



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 13 ■ PART 3 ■ OCTOBER 1994

- Guest Editorial: The Need for an Architectural Approach** 185
Andy Valdar

Architecture

- Building on the Architectural Framework** 186
Hill Stewart

- A New Structure for London's Public Switched Telephone Network** 192
Simon Wood and Roger Winterton

Human Factors

- User-Friendly Weapons for the Competitive Fight** 201
Mike Atyeo and Rob Green

Telecommunications in the 21st Century

- New Network Infrastructures for the 21st Century** 206
Part 2—Ultra-fast self-routing networks
David Cotter, Ian Hawker, Alan Hill and Vivek Tandon

- The Future of Learning** 213
Professor Alistair G. J. MacFarlane

- Distributed Management of Future Global Multi-Service Networks** 221
Robin Smith and David G. Griffiths

The Global Network

- Planning the Concert Packet Services Network** 227
Grace Hwang and Paul Seah

- Ensuring Quality in Design of Customer Networks** 233
Paul Warren

- The Design and Implementation of the BOXER Network** 240
Alan Pritchard

- The 'Value' of the Telecommunications Engineer in Eastern and Western Europe** 246
David Corrie

- 33rd European Telecommunications Congress** 250
David Corrie

- Journal Awards for Volume 12** 254

Regular Features

- Field Focus** 256

- BT News** 257

- Book Review** 260

- Industry News** 261

- IBTE Contacts** 264

Board of Editors

Chairman John Prior, M.Sc.

Managing Editor Paul Nichols, B.Sc. **Secretary** Stephen Wilson, B.Sc. **Treasurer** Chris Fennell, B.A., F.C.M.A.

Alan Bealby, M.Sc., F.I.E.E.; Richard Beet, M.A.I.E.(DIP.), F.R.A.S., F.B.I.S.; Martin Evans, B.Sc., Ph.D.; Warwick Flury; Fred Landery, B.A.;

Sam Quinn, B.A., DIP.M.; Murray Trantor, B.Sc., C.ENG., F.I.E.E., D.M.S., M.B.I.M.; Keith Ward, C.ENG., F.I.E.E., M.B.I.M.;

Syd Webster, I.ENG., M.I.E.E.I.E., M.B.I.M.; Brian Wright, B.ENG.

Published by: The Institution of British Telecommunications Engineers, 2-12 Gresham Street, London EC2V 7AG

ISSN 0262-401X

Andy Valdar

The Need for an Architectural Approach



the greater emphasis placed on the architectural process is one of cultural change

Few would disagree with the statement that the environment in which telecommunications companies operate has changed radically over recent years. Companies that started life as domestic network operators in a monopolistic environment have found that, with the necessity to interconnect to other networks, they no longer have the degree of control over network solutions that was traditional. Reducing financial operating margins has also fuelled the desire to move into areas of activity higher up the value chain, and expansion policies have resulted in moves into the global arena. At the same time, competition has increased the need to become more efficient and to use standard solutions wherever possible.

The ever increasing variety of user requirements has led to an escalation in network complexity that has the potential of defying even the best brains. Not only does high complexity introduce its own cost penalty, but it also inhibits future network and service enhancements; hence the need for an architectural approach.

The objectives of an architecture for a network operator are to set the network structure, routing, signalling, transport, intelligence, operation and service management in a framework that meets the business requirements. While one of the aims of an architecture is the minimisation of total costs, there are other, sometimes overriding, criteria. These include the speed of bringing new services to market, the ability to interconnect appropriately with other network operators and service providers, the flexibility for supporting future unspecified services, etc.

Necessarily, the agreement of architecture is not an easy process even within the confines of one company. This arises because the deployment of an architecture effectively results in the right of the individual designer to control all the decisions leading to the delivery of service being exchanged for the greater good of improved service interworking and cost reduction as a whole. Thus the greater emphasis placed on the architectural process is one of cultural change because the control of the designer is transferred to the architect.

Given this background, the introduction of a strengthened architectural process within BT has been no less difficult than might have been expected. However, the benefits have been worth the effort necessary. In introducing architectural control, care was taken to ensure that the designer was still able to influence the outcome of the process and that it was not so rigid that the design stage became extended to the point where product delivery was delayed.

The article beginning on page 186, the first of a series of articles on architecture in the *Journal*, takes a more detailed look at some of these issues. Further articles are planned which will cover in more detail this increasingly important subject.

Andy Valdar

**Head of Network Strategy,
BT Worldwide Networks**

Building on the Architecture Framework

Architecture provides the framework within which good design can take place. This article explores the significance of architecture to a communications company and provides an overview of BT's architectural approach to the development of its products, systems and services.

Introduction

'Whenever I use a word, it means exactly what I choose it to mean, neither more nor less.' (Humpty Dumpty, *Alice Through the Looking Glass*)

How Humpty Dumpty would have loved the word 'architecture'! Asking a group of computer or telecommunications experts to define it is a guaranteed way of starting an argument. There will be at least as many definitions as there are people present, and, the closer the definitions, the more acrimonious the debate about which is correct.

The problem is that most people cannot recognise architecture as a multi-dimensional concept. They focus on their own part of the architectural process and dismiss other parts of it as activities in which no self-respecting architect should be involved. In fact the architectural process covers every aspect of design, from business modelling through to preparing procurement specifications for system components.

The Architectural Approach

The importance of architecture to a communications business has always been recognised. From the beginning of the century, the world telephone system has depended on a consistent architectural approach to ensure universal interworking, in spite of large numbers of independent network operators and equipment suppliers. In one sense, the early telephone network architects had an easy task. At the end of the nineteenth century, Bell and Strowger had defined what the system was for and, to a large extent, how it functioned.

Architecture was principally concerned with engineering issues such as selecting the cheapest or most efficient signalling protocols and transmission techniques.

The architect's life became more complex when computers became associated with telephone networks. Computers can be configured and programmed to do more or less anything that can be described in numbers. Before even thinking about what kind of machine or programming language to use, the architect had to determine what the system was for and how it related to the company's business. Was it to control an electronic telephone exchange, monitor the status of a transmission link or prepare the customers' bills? Did it carry out an activity which was key to the company's competitiveness or one which merely had to be done as part of the daily routine?

The architect has to take a holistic view of the company's operation and BT's architectural approach has evolved in this direction. A system development starts by examining the business processes to be implemented and the way that information flows through those business processes. From an understanding of that pattern, it is possible to map processes and data onto the most appropriate systems and subsystems. These are then implemented by using one of a limited number of hardware and software platforms.

This approach does not come from an academic concern for engineering tidiness. It is firmly rooted in the commercial realities faced by a modern communications company. BT must generate new revenues and drive down its cost base by automating its business processes on a highly

BT's architectural approach aims to confine the impact of business decisions to a limited number of systems or subsystems that can be upgraded rapidly and inexpensively.

flexible, yet cost-effective, set of communications and information systems. Architecture provides the mechanism for identifying and implementing the requirements for those systems.

Architecture and Business Strategy

BT, like other large organisations operating in a competitive environment, needs a complex yet flexible business strategy. The strategy must reflect BT's ambition to be a major global player, but must respond to commercial, technical and regulatory pressures. Key elements of the strategy are cost reduction, improved quality of service and a broader product and market base. To be competitive, BT needs to introduce new services quickly and offer them as integrated components of an overall portfolio. This means a family of systems fully integrated to provide:

- the appropriate linkages between customer service and delivery platforms, and
- the appropriate linkages between individual services.

To meet these objectives, system development has to be proactive, providing solutions which not only satisfy current requirements, but are 'future proof'. Business decisions, say to launch a new product or change the pricing structure for an existing one, will inevitably have an impact on BT's systems. If those systems are designed around a particular organisational structure or the characteristics of individual products, the impact will be large and widespread, resulting in expensive and time-consuming redesign. BT's architectural approach aims to confine the impact to a limited number of systems or subsystems that can be upgraded rapidly and inexpensively. This reflects the company's overall strategy of responding flexibly to the market-place.

Development of BT's Architecture

Architecture covers every aspect of system design, from determining the system's role within the business, identifying the processes and data needed to carry out the tasks, grouping these in a logical way, mapping them onto specific hardware or software components, and agreeing a set of technical standards that facilitate the procurement and guarantee the interworking of those components.

BT's architectural approach is based on *Information Engineering*, a systems design methodology developed by the James Martin organisation, and used by many leading systems developers.

It starts from the *BT business model*, which is a hierarchical catalogue of the fundamental processes required to run BT's business. The process model answers the question 'What does BT do?'

From it is derived a *logical architecture*, which describes how those business processes fit together to carry out particular tasks. The logical architecture answers the question 'How does BT do it?'

Beyond the logical architecture is the *physical and data architecture*, which maps the logical groupings of

business processes and data on to specific hardware and software components and shows how those components fit together to form systems and subsystems. The physical and data architecture answers the question 'What gets done where and how do the bits fit together?'

Underpinning all this is the *technical architecture*, which defines the technical characteristics of a preferred set of components to be used in system design. The technical architecture answers the question 'Which components and tools should be used to build this system?'

Each aspect of the architecture can be considered at various levels of detail. The top level view of the technical architecture is that BT will use open standards wherever possible. At the most detailed level are lists of specifications which, for example, show how a signal should be coded for transmission over a satellite link.

Figure 1 shows the various aspects of architecture considered within BT and indicates its central role in the design process.

Architecture also has a time dimension. The technical, commercial and regulatory environment of the year 2000 will be different from that of today and BT's systems will have to

Figure 1—Aspects of the BT architecture

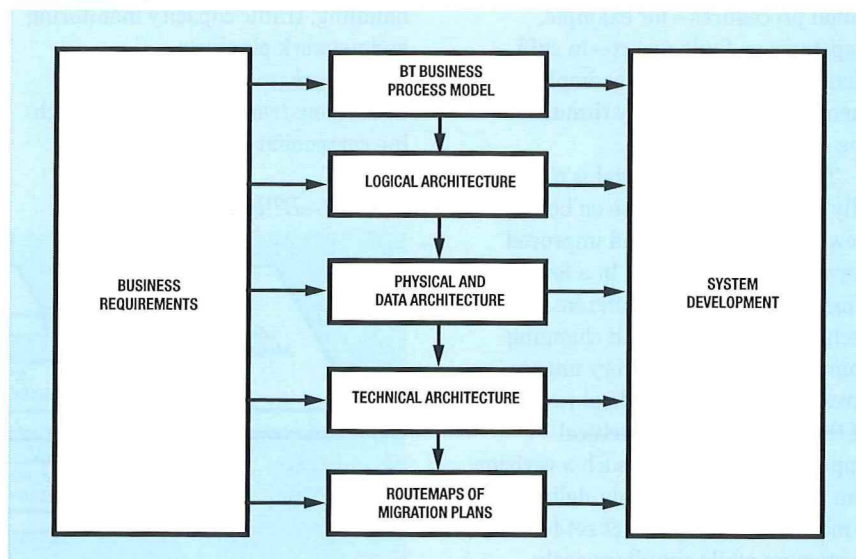


Figure 2—The BT business model

evolve to cope with it. Each aspect of the architecture must therefore include a migration plan showing how that evolution will be managed.

The following sections discuss the various aspects of the architecture in more detail and show how routemaps are used to guide the development of detailed implementation plans.

BT business model

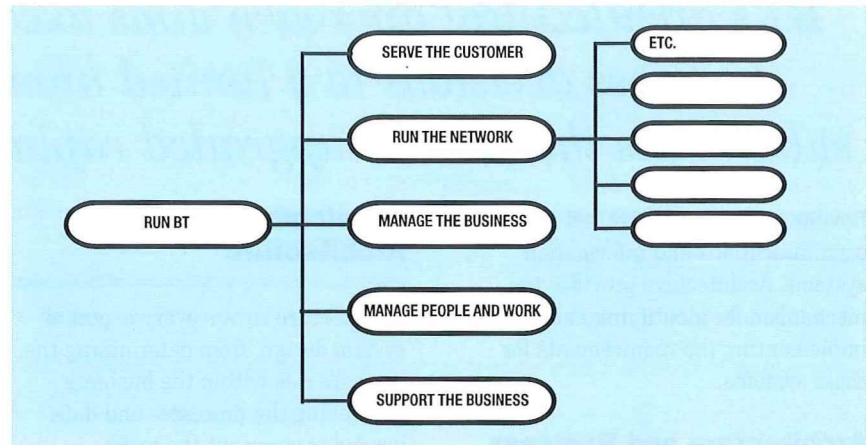
The BT business model (see Figure 2) is a hierarchical catalogue of the processes needed to run BT. It currently identifies four top-level processes which drive the company's activities:

- Serve the Customer
- Run the Network
- Manage the Business
- Manage People and Work

A fifth top-level process—Support the Business—contains non-operational activities common to all the operational areas.

These top-level processes are progressively decomposed into more-detailed sub-processes. The sub-process descriptions define the input and output data and specify what the sub-process does to that data. After perhaps six or seven levels of decomposition, they specify specific operational procedures—for example, registering a fault report—in sufficient detail for them to be implemented in the same way throughout the entire BT operation.

The BT business model is continually being refined to take on board new types of product and improved ways of doing business. In a few years time it may look very different. New technology, together with changing commercial and regulatory imperatives, may require a radical revision of the model. An architectural approach ensures that such a revision can be managed smoothly, delivering a more attractive product set to customers while simultaneously



reducing the company's operational costs.

When planning new introductions, a service or product manager maps the requirements onto the evolving business process model. The aim is to make the minimum set of changes consistent with providing customers with a consistent look and feel.

Logical architecture

The logical architecture is an implementation-independent view of a system, which presents the abstract functions (process groupings) that a system has to perform and the data items involved. BT's top-level analysis of its business and operational needs has developed a logical architecture which broadly covers the areas presented in Figure 3.

Service management covers the processes involved in providing, maintaining, monitoring and billing of services.

Network management covers the processes required to plan and maintain the BT network. This area includes processes such as fault handling, traffic capacity monitoring and network planning.

Network technology covers the underlying transmission and switching component of BT operations.

Business management covers activities such as portfolio management, preparing business plans, monitoring performance against them and providing high-level management information systems.

Resource management covers monitoring and organising BT's resources. This includes personnel, logistics, work management and maintaining the company's financial accounts.

This framework has been elaborated to a reasonable level of detail and has been successfully used as a focus for much of the planning work for BT's information systems. Examples include development and refinement of the system needs for billing, those for fault handling and the strategy adopted for customer data.

Technical architecture

As mentioned earlier, BT is committed to the use of open and, preferably, internationally agreed standards for its products, systems and services. The problem is that standards are continually evolving, some are incomplete or ambiguous and others are simply not implementable. The technical architecture therefore provides

Figure 3—BT logical architecture




Figure 4—An example of a routemap

schedules of the preferred standards, together with clarifications of the ambiguities and proprietary specifications to fill the gaps.


BT's *Network Recommendations* document the standards, and the options within them, which should be used in the construction of the company's networks. The aim is to facilitate competitive procurement, smooth interworking with other network providers and support for applications provided by other parties.

The *Technical Architecture Manual* (TAM) provides similar guidance for the company's information system designers. It recommends, for example, a limited number of hardware platforms, operating systems and databases. It also specifies a preferred set of development environments and tools for use within those environments. The aim is to keep support costs low, facilitate interworking between the company's systems and promote reuse of components between systems.

Physical systems and data architecture

The physical systems and data architecture is derived from the logical architecture, within the constraints of the technical architecture. At its highest level it defines a small number of application areas where systems are grouped to carry out the tasks defined by the logical architecture. At a more detailed level, it shows how the individual hardware and software components are configured to realise a particular system or subsystem.

BT's physical systems and data architecture are documented in a number of ways. The public network architecture is documented in the *Network Master Plan* and the *Network Evolution Roadmap*. The *Applications Architecture Manual* (AAM) describes the physical architecture and migration plans for the company's information systems. This is supplemented by 'Picture Books' which provide a



further level of detail in the areas of service management and network control. In addition, the *Data Architecture Manual* (DAM) acts as a repository for the company's data standards and provides guidance on the management of those standards and the creation of new ones.

This rather complicated set of documentation is a legacy of the multiple development routes which used to exist within the company. Work is now in hand to rationalise it and adopt a consistent way of presenting the information.

Routemaps

Routemaps provide the time dimension needed by the architecture. They chart a path between where we are now and where we want to be and identify a set of measurable events along that path where we change direction. The events identified in a routemap will generally be one of the following three types:

- functional data (for example, product enhancements, platform migrations);
- non-functional attributes (for example, performance, security, flexibility); and
- financial data (for example, event costing and cost structure).

Routemaps are best expressed as graphical diagrams of migration plans (see Figure 4 for an example), accompanied by a schedule of timed events with named owners for each. Ownership is important, as routemap development needs to be an integral part of the planning process and the

end users have to 'buy-in' to the routemap. This means that the routemaps must be realistic, balancing the interests of the various stakeholders and accepting occasional tactical detours from the 'ideal' migration path. BT has identified four distinct, but related, types of routemap:

- *Product routemaps* are product or service evolution plans which identify the product and portfolio managers' vision of how products and services should evolve during the next five years.
- *Physical design routemaps* describe the physical implementation of the product or service in terms of the hardware and software that has to be linked together to deliver them.
- *System design routemaps* describe how individual systems (for example, billing, order handling, fault handling) evolve to meet the needs of the portfolio.
- *Logical design routemaps* show how systems map onto the company's business processes.

These four types of routemap can be customised to meet the detailed needs of the many different classes of user within the company (for example, portfolio management and network evolution). They form an important link in the architectural processes which lead to a coherent product portfolio, an infrastructure to support that portfolio and a set of business processes which will help to make BT a world-class global communications provider.

Architecture is only of value if it results in the timely development of cost-effective systems that enable BT to deliver attractive products to customers.

For the information system planner, they form a mechanism for planning development and a highly communicable method for assessing progress and conformance to company objectives. As explained in the next section, BT has conducted detailed routemap and planning exercises and routinely uses routemaps as part of a conformance assessment procedure.

Implementation

Architecture is only of value if it results in the timely development of cost-effective systems that enable BT to deliver attractive products to customers. The process analysis, the production of logical models and the development of routemaps, interesting though it may be, is only a means to that end.

Purist architects would argue for a totally unified family of systems to deliver all of the company's products and services in a totally standard way. Pragmatic architects accept that their role is to minimise diversity by homing in on a manageable number of systems to implement a particular activity, ensuring that those systems run on a limited set of platforms and that they share a common, but probably not unique, communications infrastructure. In other words, pragmatic architects balance the economies of scale against the diseconomies of scope.

They also accept that systems have to evolve in response to changing business need and technological opportunities. The challenge is to balance these demands and increase stability without losing out on opportunities.

The planning process involves using the architectural framework to produce credible product and system routemaps. The events listed in the routemap are used to derive the project plan required for the development or enhancement of that particular product or system. Obviously, there is need for iteration, taking into account the resources available for development, the operational implica-

tions for roll-out and new requirements coming in from the marketplace.

To complete the loop, there is a need to assess the extent to which technical choices actually satisfy business needs, in particular the financial targets. To achieve this objective, BT has created a high-level management forum which examines the business cases for all major developments. In doing so, it is guided by a *design council*—a panel of experts who can assess the architectural implications of a proposal and ensure that it is cost effective over its full life cycle. If the proposal does not align with the overall BT architectural strategy, the proposer has either to amend it, demonstrate how the initial design will migrate towards conformance, or justify it as a one-off development meeting a key tactical need.

It should be emphasised that this mechanism is a safety net rather than a police force. The major thrust to promote the architecture is proactive, communicating the strategy to developers and working with them to ensure that, wherever possible, architecturally conformant systems are used for delivering new services.

The different components of the BT architectural strategy are all necessary to produce the benefits that can be derived from architecture. The abstract framework has to be developed in sufficient detail for system planning, otherwise the investment is wasted. Processes have to be in place to evaluate compliance with the architecture, otherwise it will not be applied as widely as it could be and the company will have a higher cost base than it should have. The architectural process aims to ensure that quality is built in and that the company can avoid wasting effort to deal with the costs of failure.

Strategic or Tactical Solution?

The architecture is not an absolute set of laws: it is a set of guidelines.

Usually, it is in the company's best interests to apply the guidelines rigidly, but sometimes a more flexible approach is needed.



The Cranfield Grid (Figure 5) is a very useful way of classifying product and system proposals and adopting a fit-for-purpose approach to their development.

STRATEGIC WILL PROVIDE COMPETITIVE ADVANTAGE CRITICAL TO FUTURE BUSINESS STRATEGY FOR EXAMPLE, CHARGECARD SERVICEVIEW	HIGH POTENTIAL LEADING EDGE, HIGHLY INNOVATIVE COULD RESULT IN COMPETITIVE ADVANTAGE FOR EXAMPLE, VIDEOPHONES, VIRTUAL REALITY
KEY OPERATIONAL ACTIVITIES CURRENTLY CRITICAL TO SUCCESS KEEPS PACE WITH THE COMPETITION FOR EXAMPLE, PSTN, CUSTOMER HANDLING	SUPPORT NECESSARY, BUT NON-CORE FOR EXAMPLE, FAX MACHINES, STOCK CONTROL

Figure 5—The Cranfield Grid

Products or systems in the *high potential* box are highly innovative: they do not yet affect mainstream business activities or significant numbers of customers. The design solutions must be rapidly realisable and need not consider the implications for core systems. 'Stovepipe' designs may be perfectly acceptable in such cases. Many developments in this box will take the form of field trials to evaluate the technology, the marketplace or both. They are often at the leading edge of technology with a very limited choice of suppliers and with equally limited scope to change the design of the solution. In such cases, rigorous application of the architecture would be counter-productive.

Products or systems in the *strategic* box are still novel but now involve significant revenues or numbers of customers. Design solutions must meet the window of opportunity and the only way of addressing the commercial imperatives may be a design that does not conform to the architecture. The



implications of such a design for other products and systems must be fully understood and there must be a strategy for migrating the design towards conformance with the architecture or amending the architecture to accommodate the requirements of this design.

Products and systems in the *key operational* box are at the core of the business. They are responsible for the bulk of its revenues and profits and are the key to commercial success. Here, designs which are totally conformant to the architecture are essential. Core products and systems must embody features that enable BT to outperform competitors in those aspects of its business that it presents as its unique selling proposition (for example, quality, price or customer service).



Products and systems in the *support* box are the 'supporting cast.' In the case of products, they may not be particularly distinctive, but they are items which customers expect BT to supply as part of an overall communications solution. In the case of systems, they perform essential business functions, but ones which are common to any large organisation. The optimum design should conform to the architecture, but is likely to contain a high proportion of standard bought-in components; indeed, it should be possible to 'badge-engineer' a product or buy in a complete system.



determining the system's role within the business, identifying the processes and data needed to carry out the tasks, grouping these in a logical way, mapping them onto specific hardware or software components, and agreeing a set of technical standards that facilitate the procurement and guarantee the interworking of those components.

It would be naïve to pretend that BT's architecture is complete, up-to-date, or universally applied. It is not and it never will be. The commercial, technical and regulatory environment surrounding BT is continually changing and the architecture must evolve to match that environment. There will always be lucrative niche products that can only be delivered in non-standard ways, new products that do not fit the current architectural vision and legacy systems that use obsolete technology, yet carry out essential business activities efficiently.

The key to the successful use of architecture is balance: balance between cost and flexibility; balance between technological novelty and proven solutions; balance between architectural conformance and speed of delivery. The architectural approach should usually be quicker and cheaper, but there will occasionally be a case for tactical solutions.

However fast current problems are resolved, technology or the marketplace will come up with new challenges and architects will have to respond to them.

Biography



Hill Stewart
BT Development and
Procurement

Hill Stewart received B.Sc. and Ph.D. degrees from the Queen's University of Belfast in 1966 and 1972 respectively. He joined the Post Office research laboratories in 1971 to work on optical communications research. From there he moved on to a range of technology forecasting and strategic planning roles before joining the Group Architect's Office in 1988 to manage the Corporate Architecture and Standards Research Programme. He is currently a Senior Advisor on Architecture in Group Systems Engineering.

Conclusion

Architecture is now an accepted concept within BT but it operates within an environment which continually questions every aspect of the business. BT is not interested in technological novelty *per se*; it is interested in the business benefits offered by new technologies. To be successful, the architecture must deliver commercial advantage.

BT's architectural approach covers every aspect of system design, from

A New Structure for London's Public Switched Telephone Network

Optimising the structure of the public switched telephone network to provide the required performance and functionality at lowest practical cost has long been a key goal. In the past, networks were designed to minimise capital expenditure, but in recent years the emphasis has changed. This article describes a new, radically simplified structure for the London network, which promises significant savings in operating costs and paves the way for the future.

Introduction

Mention the London telephone network and it will conjure up different images to different people, but for many, even vaguely connected with BT's London operation, words such as 'complexity' and 'scale' will undoubtedly come to mind. While growth, and hence scale, is to be desired, complexity creates its own additional costs and reduces the manageability of the network.

In 1991, a team within BT's London Planning Organisation conceived the idea of a target network structure. This target network vision was to guide all further developments of the network towards a simpler lower-cost structure. Although there were many aspects to the target structure, one facet of the vision was a simpler routing arrangement for public switched telephone network (PSTN) calls carried across the network. After extensive study during 1993 and the preparation of detailed costings, a decision was made to implement a radical restructuring of London's PSTN routing plan in a scheme which is now known as the *London simplified network structure* (LSNS). This article describes the background to, and nature of, the London Simplified Network Project.

Network Structure Principles

In the case of London, and a few other major cities, the size of the population and the geographic spread of the built-up area has always exceeded the

capacity and reach of any single exchange. Figure 1 shows the geographic boundaries of the current London administrative zone. Therefore, unlike many provincial towns and villages which have natural geographic boundaries, multiple exchanges are required to serve the one locality.

Routing between such exchanges becomes a rapidly more complex issue as the number of exchanges increases. Theoretically, the number of inter-exchange junction routes required for each exchange is $N(N-1)$ where N is the total number of exchanges (routes assumed to be unidirectional). This is the extreme case of a fully-interconnected (meshed) network. In the case of the pre-modernised London network, circa 1980, there were approximately 500 discrete exchanges in some 250 buildings, leading to a theoretical requirement for 249 500 inter-exchange routes, each exchange having a potential requirement for a total of 998 incoming and outgoing routes to other exchanges (even before trunk and international traffic has been considered).

This requirement to fully mesh the exchanges is illustrated in Figure 2. An alternative scenario is to connect all exchanges to one central tandem exchange. There are various practical and economic problems with both these arrangements. Fully meshing the exchanges requires mandatory *direct* routes between all units regardless of the anticipated level of traffic even where very little community of interest exists as, for example, between the northern and southern




Figure 1—BT London administrative boundary

suburbs of London. Such routes would be grossly under-utilised leading to a poor return for the investment incurred. At the other extreme, with all traffic routed via a central tandem exchange, there is not only the purchase of an expensive item of intermediate switching equipment (the tandem), but also considerable investment in the transmission plant required for transferring the traffic from any particular neighbourhood to and from the central switch. In any

event, in the case of London, total traffic demand far exceeds the capacity of any single tandem exchange ever built.

Practical Network Arrangements

Since the very early days of the London telephone service, a compromise between the 'all-tandem' solution and 'all-direct-route' solution has been considered necessary and economic.

Local exchanges which had a high incidence of traffic between them were provided with direct routes; exchanges with little traffic between them used a tandem routing arrangement. Complex rules were drawn up to determine whether tandem or direct routes should be provided in order to maximise the throughput of calls for a minimum level of capital investment. Given the capital constraints that London telecommunications traditionally experienced, there was a strong bias towards providing direct routes, expenditure on tandems not providing any obvious revenue benefits. From the 1920s to the 1970s, a highly complex arrangement of tandems, sub-tandems and sector tandems developed in line with economic and growth pressures. Figure 3 illustrates routing as it occurred circa 1980. For an overview of the tandem configuration and routing plan, readers are referred to the April 1974 edition of the *Journal*¹.

Development of the Tandem Network During Modernisation

At the start of the digital modernisation programme in the mid-1980s, there were some 40 analogue tandem

Figure 2—Tandem and direct-route arrangements for local-call routing

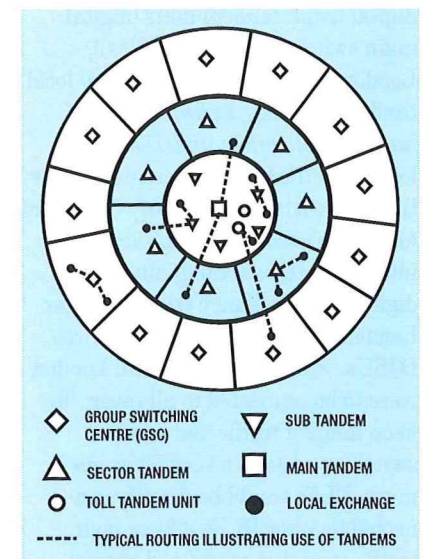



Figure 3—Pre-modernisation tandem structure for London (circa 1980)

Figure 4 – 1984 digital routing plan

and trunk switches in London, but still around half of all local traffic was carried on direct routes between exchanges. It was in this period that the current inter-exchange digital routing plan was devised for the modernised network and, with the heavy capital investment required for local switches, it was not surprising that there was again a bias towards the use of direct routes rather than tandem switches.

An issue preoccupying planners' minds at this time was the difficulty of transition between an all-analogue network and an all-digital one and the need to bridge the two while modernisation was in progress. In the early days of the introduction of System X, therefore, some exchanges were opened with up to 250 discrete routes linking them to both the analogue network and the new digital one. This large number of routes was found to cause major problems for design, dimensioning, building and subsequent administration, and from these concerns sprang the beginnings of the current simplified proposal.

The 1984 Digital Routing Plan

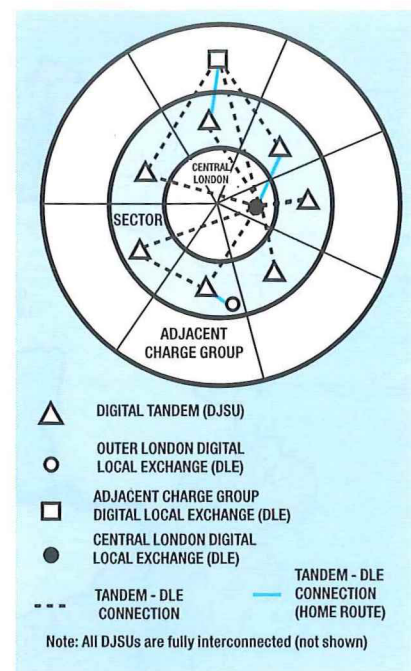
Notwithstanding the temporary problems of transition, the 1984 digital routing plan was, in fact, fairly straightforward. In addition to the digital trunk tandem units (digital main switching units (DMSUs)), London was to have seven digital local tandem switches, known as *digital junction switching units* (DJSUs) located in the buildings created under the 1965 sectorisation plan². (The term *tandem* is used to describe any intermediate switching unit.) Each digital local exchange (DLE) in outer London would be connected to three DJSUs, while those in central London were to be connected to all seven. To keep tandem traffic and hence switching costs to a sensible minimum, DLEs would be directly connected to other DLEs if busy-hour traffic exceeded predefined critical

traffic values for those routes, based on capital-cost criteria. To avoid the proliferation of very small routes, a minimum direct route threshold of 6 erlangs was established.

To provide resilience, DLEs would have up to three routing choices for any particular destination. These choices are provided at the local exchange by using the automatic rerouting (ARR) or automatic alternative routing (AAR) facilities. For each DLE, one DJSU was nominated as a 'home' DJSU. A DLE originating a call would route that call, if possible, directly to the home DJSU of the destination exchange. If the call failed to get through, the originating DLE would make second and third attempts reroute the call via an adjacent DJSU. Only if the third attempt was unsuccessful would the call fail. Directly-routed traffic is also normally allowed to alternatively route via the tandem network except where the direct route is very large, in which case it has its own inbuilt resilience by using diverse transmission systems. Figure 4 illustrates the '1984' routing arrangements.

All routes are dimensioned to give an acceptable grade of service to first-choice traffic but, as shall be noted later, this routing arrangement suffers from the disadvantage that some first-choice routes become very much larger than others used for second-choice traffic, leading to a less-resilient network than might otherwise be the case.

During the late-1980s, emphasis in BT was rapidly swinging away from keeping capital costs low to reducing ongoing current account costs. Clearly, such a vast and complex network of routes requires a huge overhead to administer and keep within an acceptable grade of service. An explosive growth in traffic and connections in the economic boom of the late-1980s almost defeated the company's attempts to keep ahead of demand and provided further impetus for the current radical rethink of network routing principles.



Drivers and Enablers for Change

Around 1989, it became clear that the sheer number of routes emanating from a typical digital exchange was causing major difficulties. Equipment costs were falling while administration costs continued to creep up with inflation. Significantly, with the advent of optical-fibre systems, transmission costs became virtually distance independent, at least within the context of London, although to gain maximum advantage from the technology available, traffic needed to be aggregated into 140 or 565 Mbit/s blocks. This favoured a few large traffic routes rather than the many small direct routes traditionally used.

As modernisation progressed, concerns were also increasingly expressed regarding the end-to-end resilience of the network. Implementation of the 1984 routing plan had resulted in greater differences between DLE-DJSU route sizes than theory had predicted, with the home route sometimes up to six times larger than each of the two adjacent (alternative) routes. Thus, under failure conditions of the home route, the grade of service deteriorated severely. A new network structure was needed to even out route sizes if resilience was to be improved.

During the rapid traffic growth of the late-1980s, it was decided to expand London's trunk and tandem network by installing a further five digital main switches in central

London. These included a new DMSU and four new multipurpose digital switching units (DSUs); intended mainly to switch traffic between central London and the home counties. In the event, predicted traffic levels did not materialise leading to excessive spare capacity. However, the exercise illustrated how difficult and expensive it is to insert a new trunk or tandem switch into a traditional fixed-routing-plan network. Large amounts of network require to be rebuilt with the accompanying rework of exchange data. It was thus clear that any new tandem routing scheme should aim to reduce the structural problems of adding extra capacity.

One further development also prompted a rethink of routing principles. This was the hitherto unused facility of the latest DLE software upgrades to distribute calls proportionately across a number of routes. The proportional traffic distribution function (PTDF) held out the prospect of being able to build a network which was much simpler to dimension and which would be more resilient to failure as a single logical route could be created from several physical routes.

In parallel with these developments in London, thinking nationally was also developing towards making a case for a simpler and more uniform PSTN structure across the country as a whole. It was recognised that this would have particular benefits for the overall network traffic management

undertaken by the Central Operations Unit (COU), in Oswestry. Thus, in 1992, a working party was formed comprising representatives from London, the COU, the structures policy unit and some zone operations units. Work progressed under the guidance of this forum to define a simple uniform national structure using existing DMSUs, but eliminating all direct DLE-DLE routes.

Call Routing via the Simplified Network Structure

The principal idea for the simplified network structure (SNS) was to connect all DLEs to just three tandems with three routes of equal size, see Figure 5. To establish a call, DLE X examines the first few digits dialled by the customer. These determine whether the call is an own-exchange or a non-own-exchange call. If the latter, the call is offered sequentially to routes 1, 2 and 3 until a free circuit is found. It should be noted that this does not represent a first, second and third routing choice since subsequent calls may be offered to routes 2, 3, 1 or 3, 1, 2. By using the PTDF facility, calls are thus distributed evenly on the outgoing routes to tandem switches A, B and C. At the tandem switch, the dialled digits are again examined to determine the destination for the call. In the case illustrated, the destination DLE, DLE 'Y' is triple parented

onto tandems D, E and F. Within the tandems A, B and C, the national number code for DLE 'Y' will be understood to require a routing to the tandem group D-E-F. Having determined the destination tandem group, switch A, B or C will distribute the traffic evenly on each of the three appropriate outgoing routes. At the destination tandem, D, E or F, the automatic alternative routing facility will be enabled to provide alternative paths to complete the call if the direct route to the destination DLE Y were busy. For example, if the route from D to Y were fully occupied, then D would attempt to complete the call via E or F.


Benefits of the Simplified Network Structure

Such an arrangement has a number of distinct advantages. Firstly, the originating exchange only requires to recognise whether or not a call is an own-exchange call or not. This means that there is no longer a need to maintain a complete set of data describing the routes to all destination exchanges. In the context of London, where new all-figure numbers (AFNs) may be introduced at up to 10 per month, this represents a significant saving of effort across over 100 DLE processors. In addition, any DLE only has three outgoing routes, each of which carries one third of the traffic. Thus, network monitoring and dimensioning is considerably reduced.

Each of the three routes will typically be very large, leading to very economic utilisation of modern transmission systems, although one drawback of this arrangement is that where there are several DLE processors in one building, traffic which was previously carried internally between processors on 2 Mbit/s ties now has to be carried to the tandem sites and back. In practice, it is found that this cost penalty is more than outweighed by other benefits.

One of the advantages of having just a few large traffic routes is that

Figure 5—Simplified network structure routing principles



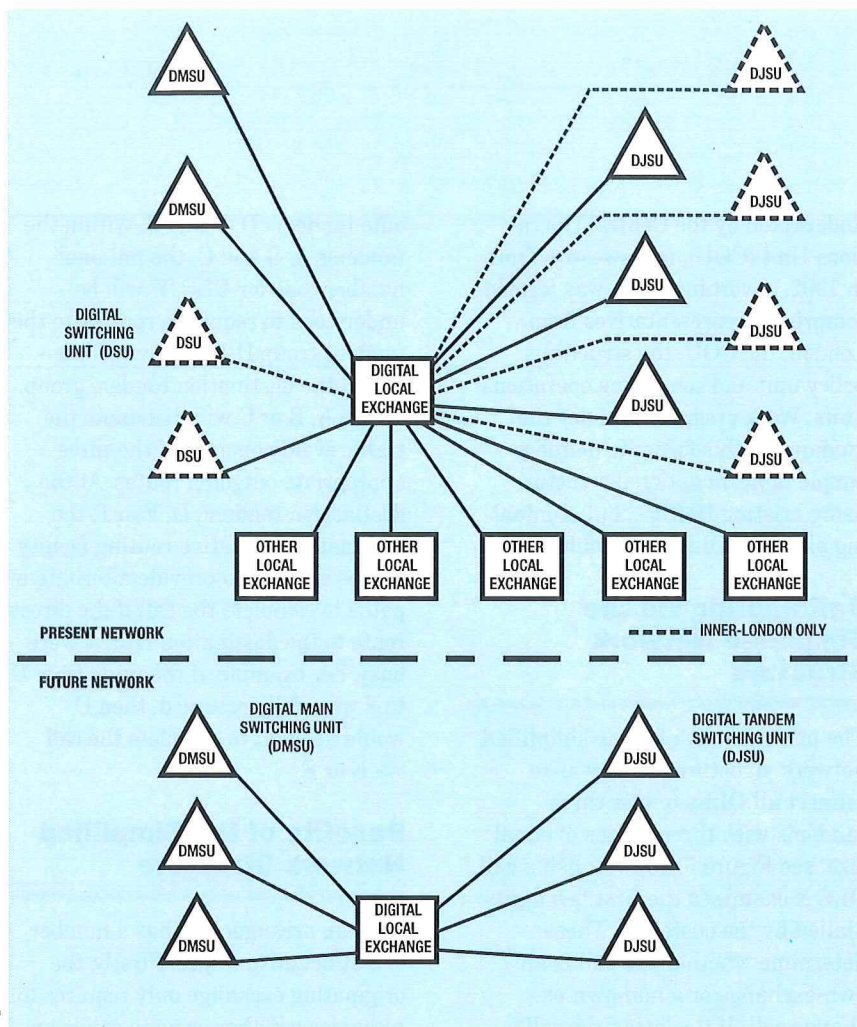


Figure 6—Impact of LSNS on digital local exchange network

will be connected to three trunk switches (DMSUs), and they will use PTFD to distribute outgoing calls evenly between the three routes. However, a separate London-only local tandem switching layer will be retained using DJSUs, as this provides useful economies in terms of the costs of change.

A large proportion of the traffic generated by London customers terminates within the London area. In fact, 11 DJSUs are needed to meet the expected local tandem switching requirements at the year 2002, while only nine DMSUs will be required for trunk traffic. The London DMSUs will need to be fully interconnected with all DMSUs nationally while the DJSUs need only be fully interconnected within London (that is, just ten routes each). Retaining separate trunk and local tandem switching functions will also reduce the amount of double switching of traffic between London destinations, thus saving tandem switch capacity.

The separate local tandem layer will be configured in the same simple structure as the trunk layer, with all DLEs connected to three DJSUs, and PTFD used to distribute calls evenly across the three routes. No direct routes will be provided between DLEs, and their standard six-route structure will represent a dramatic simplification by comparison with existing arrangements—see Figure 6.

In creating the new structure, it is planned to use the existing tandem switches *in situ*, albeit with certain changes in function. The existing seven outer London DMSUs and DJSUs will continue to perform their trunk and local tandem functions respectively, but in inner London only two DMSUs will remain. Of the other inner London tandem switches, one DMSU and three DSUs will become DJSUs, while one DMSU and one DSU will be offloaded and eventually closed.

The new switching structure and physical switch locations are illustrated in Figure 7.

The LSNS retains the benefits of simplified routing data in the DLE,

proportionately less capacity needs to be installed in order to deliver the same grade of service. Table 1 shows the efficiencies achieved in carrying 200 erlangs of traffic over three routes as opposed to 5, 10, 15 or 20 routes.

Table 1 Route/Circuit Requirements

Number of Routes	Circuits Required to carry 200 Erlangs*
3	$3 \times 82 = 246$
5	$5 \times 52 = 260$
10	$10 \times 29 = 290$
15	$15 \times 22 = 330$
20	$20 \times 18 = 360$

* This is the number of circuits required to carry 200 erlangs with a grade of service of 0.013; that is, giving a probability of successfully routing a busy hour call of 98.7% (without alternative routing).

In addition to the benefit obtained by having large routes, the three physical routes effectively form one logical route, allowing a further small reduction in the capacity required while still maintaining the grade of service.

While the above describes an ideal arrangement, practical considerations, in particular the costs of changing from one routing scheme to another, were found to require some compromises. Unfortunately, nationally, it has not yet been possible to justify the costs of change by the anticipated reduced running costs. In London, however, a scheme based on the SNS principles has been devised which shows a cost-benefit within three years.

London Simplified Network Structure

The London simplified network structure (LSNS) builds upon the basic principles developed for the SNS. That is to say, all London DLEs

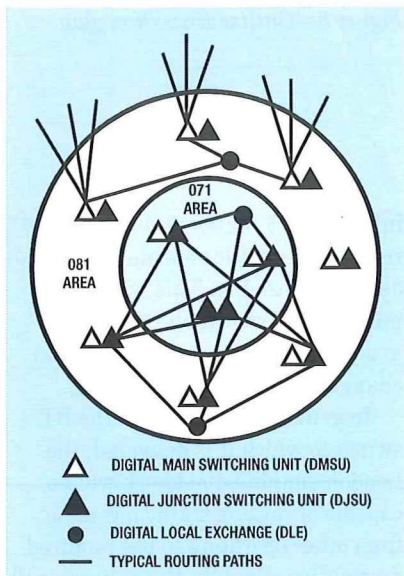


Figure 7—New London tandem switch structure

as the latter still has a very simple decision to make for each dialled call—is it an own-exchange call or should it be routed to a DMSU or to a DJSU. To get the full benefit of this simplicity, unused exchange codes (London AFN or national number group (NNG) DE digits) will not be identified by the DLE. All codes will be routed to DMSU or DJSU as appropriate, and that switch will be responsible for recognising the unused code and ensuring that the correct message is returned. Thus, for most calls, the DLE will need to examine only the first 3 or 4 digits (for example, 071, 0923, 0225) of the dialled national number*, in order to be able to determine the appropriate routing.

A further novel feature of the LSNS is the ability to dispense with geographical catchment areas for the tandem switches. The London area, although heavily populated with telephone users, is relatively small in geographic terms, and, with the advantages of modern transmission technology, means that route distance from DLE to tandem within London is not a significant factor in overall network costs. Thus, any DLE can be parented on any three tandems, and it is therefore possible to balance tandem loadings evenly. This capability will be limited by the need to minimise the number of groups of

three tandems used for DLE parenting, in order to reduce the amount of routing information to be held by DMSUs and DJSUs—each set of three parents requires to be separately identified in data. But in practice, balanced tandem loadings should be readily attainable. This will confer benefits in terms of switch capacity utilisation, and resilience (provided that sufficient spare capacity is maintained to ensure headroom under switch failure conditions).

Implementation Methodology

The development of a suitable, simple PSTN structure to meet London's future represents a significant milestone. However, before such a structure could be introduced, it was necessary to develop an implementation scheme whose cost could be readily recouped by the savings created. As has been explained previously, the lack of a cost-effective implementation scheme has so far presented an insurmountable barrier to the early introduction of SNS throughout the rest of the UK.

The solution to this problem in London is to be found partly in the choice of physical structure and partly in the implementation methodology. It has already been explained that the existing DJSU/DMSU/DSU structure will be converted to a simpler DJSU/DMSU structure, by using existing switches in existing locations but with changes of function for four of the inner-London main switches. Thus, implementation involves no expenditure on switching equipment.

Similarly, by careful choice of parents for each DLE, it will be possible to make use of the great majority of existing DLE-to-tandem routes. Thus, while the introduction of the third DMSU parent for DLEs in outer London, and parenting changes in inner London, will create some new route requirements, expenditure on new transmission equipment will be minimised. It

should be noted that while this approach will tend to maintain the old geographic catchment areas for tandems at the outset, there will be no continuing geographic constraint once the scheme is in operation.

Further savings in capital costs will be achieved by coordinating the implementation of the LSNS with the TXE4 replacement programme. London's remaining TXE4 electronic exchanges are all programmed for replacement by 1997, and it is therefore desirable to minimise LSNS implementation expenditure on TXE4 networks. Those TXE4s replaced before or during LSNS implementation will require virtually no work to be done. Their replacement processors will be provided with the appropriate LSNS network connections—at lower cost than if old-style network connections had been provided. Some simplification of routing data will be carried out at the 30 or so TXE4s which will remain after the introduction of LSNS, but again no new network provision is envisaged.


As a result of these factors, the overall cost of implementation will be relatively small, with expenditure on equipment outweighed several times over by manpower costs. In fact, the principal resource requirements will be for rebuilding the exchange data (necessary at all London digital switches) and for circuit provision. However, such is the cost effectiveness of this approach, that implementation is expected to show a positive return in less than three years.

Practical Transition Plan

The methodology described opens the way to a straightforward two-stage transition (see Figure 8). In phase 1, the new trunk switching layer will be established as follows. First, the necessary new routes and augmentation to existing trunk routes will be put in place. Then new data will be introduced at DMSUs nationally, and the new routing plan implemented for traffic incoming to London. Finally,

* After the national code change, in April 1995, an extra digit will need to be checked as codes will become 0171, 01923, 01225 etc.

Figure 8—Outline transition plan



new data will be introduced at London DLEs, and the new routing plan implemented for outgoing trunk traffic.

Once the new trunk switching layer has been introduced, there will be only two inner-London trunk tandems. The four DSUs and the two other DMSUs will have been offloaded to open the way for the implementation of phase 2, when four of the offloaded switches will be reused as DJSUs, while the remaining two units will eventually be taken out of service.

Implementation of the local tandem layer in phase 2 will follow a similar pattern to phase 1. That is to say, the first step will be to provide any new routes and augment existing routes where necessary. Then the new DJSU data will be introduced, and the new routing plan will be implemented between DJSUs. Finally, new simplified data will be introduced at DLEs to effect the new traffic routing plan for outgoing traffic.

Implementation of this transition plan has already begun, and is expected to take around two years to complete.

Implications for International and Interconnect Traffic

Currently, international traffic to and from BT's London customers is routed via DJSUs in outer London and via

DMSUs in inner London to the international gateway exchanges. Until recently, some large London DLEs even had direct routes to the international gateways. Once the new network structure is introduced, all international traffic will be routed via the London DMSUs—a further useful measure of standardisation.

The extent of interconnection between BT's PSTN and the networks of other licensed operators (OLOs) has increased rapidly during the past few years. Now, in addition to interconnect with Mercury Communications Ltd., Cellnet and Vodafone, there are requirements to connect with new national network companies, new cellphone operators and local cable-TV companies, as well as with some other operators with more specialised service offerings. One end of each such interconnection must terminate on a BT switch—usually at a DMSU or DJSU exchange, but occasionally at a local exchange. Thus, in undertaking major changes to the BT network switching structure and traffic routing plan, very careful consideration must be given to the impact upon interconnect arrangements.

In determining how to best deal with the interconnect implications of introducing LSNS, one fundamental principle has been established: to retain existing

interconnect routes and their recognised traffic routing function as far as possible. This will minimise the impact upon OLOs and avoid unnecessary expenditure on change to interconnect routing.

In general, regardless of the BT switch to which it is delivered, the London simplified network will be capable of routing traffic incoming from other operators to the required destination. Equally, the network will be able to direct traffic outgoing to OLOs to the appropriate point of connection.

The main changes to interconnect routing will occur in consequence of the closure of a DMSU and DSU. In practice, this will result in significant changes for only one operator. Discussions are already in hand to develop plans for a phased transfer of that operator's traffic onto retained DMSUs. Discussions will also take place with other operators to ensure that all are fully aware of BT's intentions and to resolve any network interconnect issues which may arise.

For the longer term, and as the requirements of other operators evolve, then routing arrangements within the BT network will also evolve, to provide the best match between the respective requirements. In this context, BT's aim would be to introduce proportional traffic distribution and balance the loading on interconnect routes. Also, the aim would be to ensure that OLO routes are connected to a BT switch with the appropriate function. For instance, interconnect traffic which is effectively local (that is, originating and terminating with London customers) would ideally route via the DJSU layer. Similarly, interconnect traffic which is effectively trunk (that is, originated or terminated with customers outside London) would ideally route via DMSUs. Such arrangements will minimise the impact upon the BT network of changes in traffic patterns due to customers migrating between BT and OLOs.

Figure 9—Adding an additional tandem in a simplified network structure

Key Benefits of the London Simplified Network Structure

The key benefits of the LSNS lie in easier management and reduced manpower requirements, leading to reduced costs for planning, building and operating the network. These benefits will flow from its twin features of simplicity and standardisation.

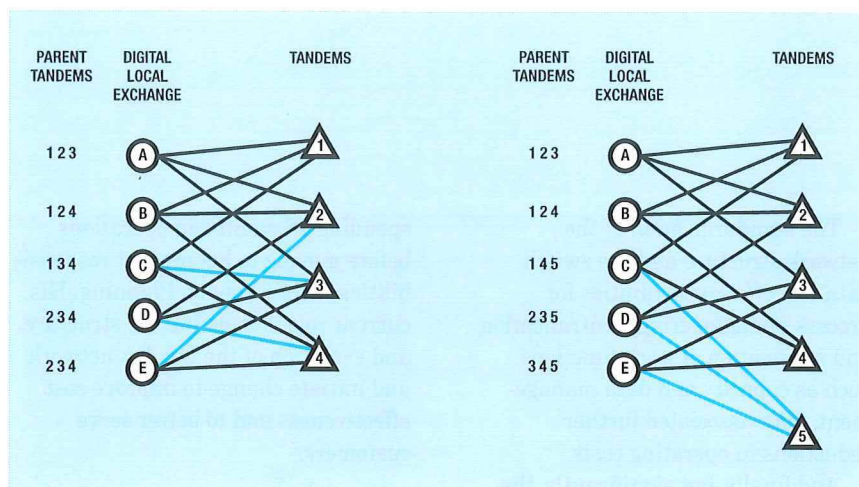
The inherent simplicity of the routing arrangements will reduce route dimensioning to a trivial task. Similarly, the reduction in the numbers of physical and logical traffic routes will simplify capacity management, while the economies of scale arising from the concentration of traffic onto fewer routes will reduce the number of 2 Mbit/s systems in the network and thus enable savings in maintenance effort (see Table 2).

Table 2 Reduction in Network Size

Network Element	Reduction Achieved
Physical routes	50%
Logical routes	80%
2 Mbit/s systems	20%

The reduction in numbers of physical traffic routes will also offer economies for BT's CCITT No. 7 signalling upgrade programme, which is moving forward in a similar time-scale to the LSNS project. The upgrade requires changes to signalling cards associated with each traffic route, and thus fewer routes means fewer cards to modify or replace.

The introduction of LSNS will result in the provision of standard data build and network configurations for DMSUs and DJSUs (eliminating the anomalous DSUs). It will also provide a significant simplification of DLE data structures and result in a dramatic reduction to the number of changes required to DLE



In the hypothetical example shown, only three DLE route changes are required to introduce the new tandem. These can be chosen independently of the locations of DLEs and tandems. In previous network structures, far more extensive changes would have been required.

data. For most calls, DLEs will no longer need to recognise the specific destination, but will simply decide whether to route via a DMSU or a DJSU. Thus, network changes such as remote concentrator unit reparenting or new processor opening will be able to take place without modification to data at other DLEs.

The introduction of an additional tandem switch would also be greatly simplified by the combination of the standard, balanced structure and the absence of geographic catchment areas. The method for introducing a new switch is illustrated in principle in Figure 9. This methodology could equally well be used in reverse to take out capacity in an environment of shrinking demand.

The Future of the Public Switched Telephone Network

No article on the PSTN structure would be complete without a look to the future of telephony. From a users' point of view, the next few years will see a rapid increase in the functionality of the switched network, with new intelligent network features added to the basic platform, and with digital transmission at customer premises becoming the norm (through the integrated services digital network (ISDN)). The simpler structure described here reduces the cost of those changes and in many ways makes them more straightforward to implement.

BT's digital PSTN platform will be in service for the foreseeable future, but overlay and alternative service platforms will take an increasing proportion of traffic. For example, BT's fast growing centrex and virtual private network services are currently provided on an overlay platform. Interconnection of the various platforms is made considerably more straightforward with LSNS.

Looking further out, a public asynchronous transfer mode (ATM) platform will be required if broadband switched services (for example, B-ISDN) are to be provided to customers. Interworking issues similar to those faced in the transition from analogue to digital technology will be greatly reduced if the narrowband platform is structured in accordance with SNS principles.

Conclusions

With its greatly enhanced flexibility to accommodate change, the new London structure should provide an excellent platform for the capital's PSTN service into the next century. Network changes will generally be easier to implement, and the introduction or withdrawal of a tandem switch to meet changing traffic demand will be much more straightforward than with the earlier network structure.

The balanced loading to tandem routes will facilitate high network utilisation, and resilience will be improved as maximum access is available to alternative route capacity.

The standardisation of the network structure and the switch data will offer opportunities for process re-engineering, centralisation and automation of work functions such as capacity and data management, with associated further reductions in operating costs.

And finally, but significantly, the simplification of the basic narrowband platform will greatly ease any future transition to broadband and intelligent networks.

Acknowledgements

The assistance of David Bye in checking various drafts of this article is gratefully acknowledged, as is his work to develop the LSNS transition plan detailed herein.

References

- 1 WHERRY, A. B. and BIRT, J. F. The London Sector Plan: Background and General Principles. *Post Off. Electr. Eng. J.*, **67**, Apr. 1974, p. 3.
- 2 DAVIES, E. and FRANCIS, H. E. Decentralization of London's Trunk and Tandem Telephone Exchanges. *Post Off. Electr. Eng. J.*, **59**, July 1966, p. 204.

Biographies



Simon Wood
BT Worldwide
Networks

Simon Wood joined BT, in 1980, after graduating from Cambridge in Engineering and following a short spell in teaching. He worked for three years on private-circuit engineering standards and provision processes before being sponsored to take an MBA at Henley Management College. He moved to work in the London Network Organisation, in 1987,

spending 18 months in operations before moving to his present responsibilities within London Planning. His current role is to define the structure and evolution of the London network and initiate change to improve cost effectiveness and to better serve customers.



Roger Winterton
BT Worldwide
Networks

Roger Winterton is Head of Network Architecture in Worldwide Networks, London. He joined the Post Office as an open entrant in 1968, working on computer-hardware appraisal and acceptance testing in the Telecommunications Development Department at Headquarters. After taking an honours degree at the City University, he was appointed Executive Engineer, in 1975, and moved to the London Regional Headquarters. He joined the Critical Path Methods Group, working on major building projects, such as Burne House, before moving into Trunk Long Term Planning. After a spell at BT Headquarters, he moved back to the London Regional Headquarters, in 1984, to become Head of the Junction Network Strategy and Tandem Planning Group, taking on his present role in the major reorganisation of 1991. He is a Chartered Engineer and a Member of the Institution of Electrical Engineers.

Glossary

- AAR** Automatic alternative routing.
- AFN** All figure number. The term AFN is used to describe the first three digits of a 7-digit customer number as found in the major cities of the UK. These digits identify the local exchange on which the customer number is to be found.
- ARR** Automatic rerouting.
- ATM** Asynchronous transfer mode.
- COU** Central Operations Unit (Oswestry).
- DJSU** Digital junction switching unit—a digital local tandem exchange found only in London.
- DLE** Digital local exchange.
- DMSU** Digital main switching unit—describes BT's main digital trunk exchanges.
- DSU** Digital switching unit—this term has been applied to the four multipurpose exchanges installed in central London which have both DJSU and DMSU functions.
- ISDN** Integrated services digital network. All of BT's inland digital PSTN is capable of carrying ISDN traffic.
- LSNS** London simplified network structure—the London variant of the simplified network structure having separate trunk and tandem layers.
- PTDF** Proportional traffic distribution function.
- OLO** Other licensed operator; that is, non-BT telecommunications operator.
- PSTN** Public switched telephone network.
- SNS** Simplified network structure—the simplified system of routing using just one layer of tandem switching to interconnect local exchanges.
- TXE4** Electronic processor-controlled local telephone exchange using analogue (reed-relay) switching.

User-Friendly Weapons for the Competitive Fight

Customers say they want products and services that are easy to use—and there is good evidence that user-friendly systems deliver more benefit to both users and suppliers.

However, selling user friendliness is not straightforward as it involves the subtleties of gaining competitive advantage from usability.

Introduction

Customers are looking for—and increasingly choosing—products and services that are easy to learn and that can be used quickly, accurately and with little effort. Reviews in the technical and popular press guide customers in their choice of service or supplier and these usually include ease-of-use ratings. This burgeoning interest in usability stems from increases in the complexity of technology, the growing sophistication of customers and the greater choice offered by competing suppliers. In response, many companies advertise 'user-friendly' products and services. Sometimes however this is no more than advertising hype, despite the fact that deceiving customers leads to disappointment and problems. On the other hand, some companies—including BT—are committed to making their products and services genuinely easy to use to gain the advantages of meeting their customers' real needs.

Customers Demand 'User-Friendly' Products

As a result of experiences with complicated, difficult to operate products such as the notorious video recorder, many consumers feel let down by designers and technology-driven manufacturers. Some say that they have difficulty even **imagining** companies providing products that can be used without carefully studying lengthy, complex instructions. Market research shows that many customers expect electronic equipment to be 'unnecessarily difficult to learn', with 'arbitrary' operating procedures involving much effort. People may put up with confusing, frustrating systems when they have no alternative—at work for example—but they increasingly select user-friendly products when they do have a choice.

Consumers' buying decisions

In the rapidly changing telecommunications industry, customers may not fully realise what a wide choice of products and services is available. They may not fully understand their communications requirements or what telecommunications can do for them. They may not have a clear picture of what benefits to expect from a service or how much it should cost. In consequence, they may find it difficult to make a rational choice between competing products.

People want products and services that they can use without undue effort or a long learning period. However, it is not always easy to use this criterion when faced with an array of competing technology in a shop, a splendour of 'special offers', and a persistent salesperson (Figure 1).

Figure 1—The buying decision






Figure 2—Video-on-demand interface looks good and is easy to use

Customers cannot necessarily spot the clues from the products' knobs and dials, buttons, labels and displays that could lead them to judge which product is easiest to use—especially if they cannot try it out in advance. In this situation, they may revert to simpler criteria, such as cost or number of features. They may even reject a product because 'if it's easy to use then it must be the basic model'—with the implication that the basic model is unlikely to meet their full requirements. The challenge for designers is to produce a product that is both easy to use **and** has the high-quality look and feel that ensures it is attractive to customers (Figure 2).

After the sale—use and evaluation

After choosing a product or service, customers start to use it and form an opinion about their selection. If it is difficult to use, they may send it back, not make full use of it or decide never to choose that service or supplier again. On the other hand, with a user-friendly product, they may use it more than they expected, buy a whole set of associated products or services, recommend the supplier to their friends and decide to continue to use that supplier in the future.

People sometimes blame themselves when they cannot easily use a product and do not recognise that this is as much a design fault as mechanical failure. They do not always return a product simply because it is difficult to operate, partly because a product is

almost never completely unusable and they can retreat to using a core set of functions that more or less meet their basic needs. However, difficulty of operation restricts usage.

Sometimes, customers express satisfaction with a product, and yet their behaviour shows that they are not exploiting its full potential. For example, many customers regard their telephone as very easy to use, but they may not even be aware of the *, # and RECALL buttons that could deliver the benefits of network services (Figure 3). Lack of awareness, complex operating procedures and lack of understanding of benefits can all deter customers from taking up or using new services.

However, market research shows that customers are changing. They are beginning to recognise the role of network-based services and communications products in meeting their needs. They are also becoming more aware of the benefits of communication and of the potential that today's technology has for providing systems that are easy to learn and to use. They are increasingly demanding user-friendly products.

Marketing usability


The intangible nature of usability makes the task of turning it into a competitive advantage far from straightforward. Advertisements often treat user friendliness as a product feature. The addition of an 'easy to use' bullet point on a feature list seems equivalent to a 'now washes

whiter' slogan for soap powder, as though usability can be added to a mature product to make it more appealing. However, no one element determines ease of use. It is created by a combination of many factors such as the physical design of the product, the matching of the information displayed by the system to the users' characteristics and limitations and a user guide that clearly explains the product in terms that are familiar to the end user (Figure 4). Usability is a pervasive quality resulting from a design process that takes the user into account throughout. The technicalities of usability performance data do not make for interesting promotional campaigns, but they can provide the basis for a set of convincing messages that will help to sell the benefits of a product or service. These benefits could include reduced learning time, higher performance with less effort, reduced costs and a more pleasant experience for the user. For business users, these benefits can be expressed in financial terms.

The Business Case for Usability

Consumer buying decisions depend mainly on attitude and perception, and the benefits to individuals are primarily greater satisfaction and reduced frustration. On the other hand, organisational buying generally has a firmer basis in predicted costs

Figure 3—The home telephone—easy to use, or is it?



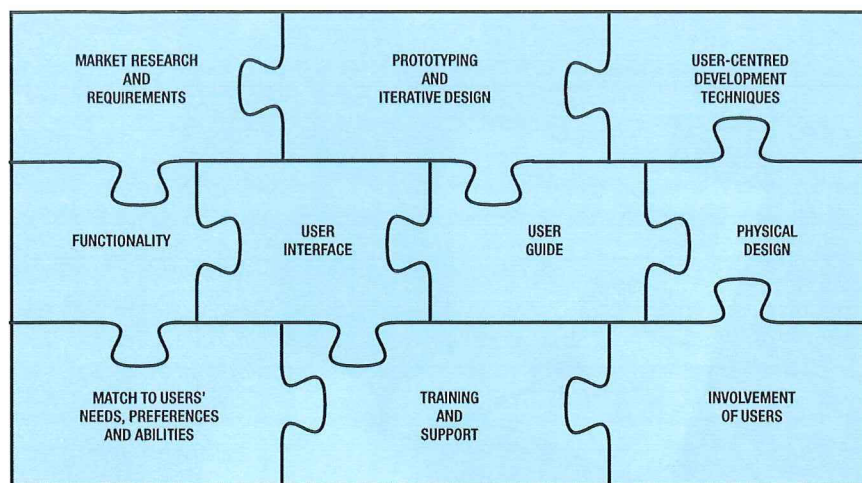


Figure 4—Some components of usability

and benefits. Financial benefits accrue to both business customers and to the companies who provide them with more usable products and services.

Benefits for customers

Attention to meeting end-user requirements sometimes has a dramatic effect. For example, a human factors evaluation before buying specialist workstations recently showed that an industry-standard product, costing only one-third of the proposed offering, was capable of supporting the users' tasks at the required performance level.

The benefits to customers result from one or more of the following:

- appropriate rather than over-engineered solution;
- decreased task time, increased productivity;

- reduced errors;
- less training;
- reduced support;
- less user disruption; and
- fewer software changes.

Companies have always been aware of the high costs of training, and of the reduction in these costs from well-designed products, from simple instruction manuals, and from on-line help. They are also becoming more aware of the distributed costs of supporting the users of poor products. In addition, businesses are more able and willing to take whole life costs and 'soft' issues, such as high staff turnover through low morale, into account.

Organisational buying decisions

Even the business market is not straightforward—and the way in which an organisation buys can represent more complexities for marketing ease of use. End users face usability problems. The buyers, decision-makers, and other people who influence the choice of a system may never have to use it. Users' difficulties may remain invisible to these people because the users find ways of working around the problems or because the difficulties and their costs are distributed invisibly among many business units. Alternatively, users' views may simply have a low priority in the decision-making process. Marketing effort is needed to educate those responsible for making the selection about the impacts of ease of use on the costs and benefits

of alternative products, systems and services.

Benefits for BT

Attention to usability also benefits service providers—by increasing market share and revenue, by reducing costs and by improving customer perceptions. For example, by making the set-up and use of network services easier, an increase in usage and revenue can result (Figure 5).

Products that are difficult for customers to use create high internal costs. The proportion of these costs due to poor usability can often be quantified by analysing helpline calls. Poor usability can result in a high rate of customer returns for which no fault can be identified. The costs of these unnecessarily returned products are also high, and include transport, storage, handling costs, internal processes for managing returns, cost of replacement or refund, time spent in negotiation with the customer, inspection and repair costs, and costs of disposal of returned goods.

Revisions to software or products may be necessary to reduce customer problems resulting in costs of distribution, revision of user guides and retraining that can rapidly erode profit margins. Supplying usable products to customers can reduce all of these costs. The costs of damaged customer perceptions, lost customers and reduced usage are more difficult to estimate, but no less significant.

Figure 5—Screenphone—one approach to user-friendly network services



'The usability of BT's portfolio has a direct impact on customer perception of BT. It has the potential of being one of the most powerful means of deriving positive product differentiation, and, more importantly, competitive advantage.'

Stafford Taylor, 1994
Managing Director,
BT Personal Communications

Figure 6—Story boards used to capture and refine design

Meeting the Demand

Companies in the 1980s producing ever more complex software applications in an extremely competitive market had to deal with the difficulties of meeting customer demands for user friendliness. The emergence and meteoric growth of the personal computer industry opened up a new market of inexperienced users. Similar increases in system complexity and range of user experience occurred in telecommunications. A number of lessons have emerged from this experience about the competencies that a company must have to deliver systems that are easy to use.

Management championship

Involving the end user throughout the development of new services requires significant changes in approach. Such changes seldom happen without senior management commitment.

'It is essential that BT's products and services are easy to use.'

Dr Alan Rudge, 1994
Managing Director,
BT Development and Procurement

Although often the subject of an explicit corporate initiative, usability may also be delivered through other avenues such as quality programmes. Modern quality programmes emphasise the importance of meeting customer needs. Ease of use is one of those needs and a user-centred development process is a key means of meeting it. In the 1992 European Quality Award, the marks were distributed among nine criteria with Customer Satisfaction clearly the highest valued, receiving 20% of the overall marks.

This route may become more important in the future, as a recent ACOST* report¹ recommended that

* Advisory Council on Science and Technology



'...the British Standards Institution should consider how to incorporate human factors and organisational design into the criteria for ISO 9000...'

In response, the BSI has recently started discussions with the British Computer Society on integrating human factors with software quality. ACOST also recommended that 'Companies should include human factors and organisational design explicitly in business plans...', as well as '...in the design and development of products, processes, and systems'.

'User-centred' development processes

Involving users throughout the development life cycle provides rapid evaluation of design decisions and inspiration for improvements². Traditionally, users are involved in initial market research and eventually in product trials but not in the intervening development stages. In user-centred design, users have explicit roles in project teams and development processes, moving away from passive roles as observers, approvers and recipients, towards active roles as co-designers, self-advocates and experts.

Usability is sometimes seen as a 'black art'—impossible to measure or manage. However, this aspect of a new product can be specified just as clearly as its market window and cost targets, and acceptance criteria can be established based on observations of real users. These criteria should at least cover the dimensions defined by

an ISO standard on usability specification³: user satisfaction, effectiveness and efficiency. A detailed cost-benefit analysis for ease of use can also be carried out, so that appropriate commercial trade-offs can be made during development.

Customers cannot necessarily express their own requirements, but are good at identifying ways of improving prototype systems. Designers often find that they cannot communicate their design in a language which users understand. Rapid prototypes allow both designers and users to get an early insight into the proposed product or service. Rapid prototyping enables early user participation, and iteration provides a cost-effective way of making appropriate design decisions at the right time, based on real user behaviour and customer needs. Even simple paper-and-pencil sketches can be used to explore both marketing and engineering issues and a growing number of prototyping tools allow more sophisticated prototypes to be developed early in the development process (Figure 6).

Business and management processes

Business processes, organisational structures, responsibilities and incentives all affect a company's success in delivering user-friendly products, systems and services. Current arrangements in BT focus marketing, design and human factors resources on ease of use when the

product manager expects usability to be critical. The continuing trends towards greater interactivity, less reliance on special technical training for users and convergence of computers, communications and other media make this increasingly important across the entire portfolio.

Conclusion

BT is in a rapidly changing and extremely competitive marketplace. Ease of use is now a clear driver in all consumer and business markets and the convergence of technologies, markets and customer demands will lead to new products and services meeting a new order of customer needs. These new products and services will become more sophisticated and integrated, but users will also demand that they are easy to use and consistent in their 'look and feel'. By implementing user-centred development processes and integrating human factors techniques into everyday business, BT will deliver more effective and efficient products, will continue to meet customers' novel needs, and provide an ever more satisfying experience.

References


- 1 Advisory Council on Science and Technology. People, Technology and Organisations: The Application of Human Factors and Organisational Design. HMSO, 1993
- 2 COOPER, MARTIN. Human Factors in Telecommunications Engineering. *Br. Telecommun. Eng.*, July 1994, **13**, p. 98.
- 3 International Organisation for Standardisation. ISO 9241. Ergonomic requirements for office work with visual display terminals. Part 11—Guidance on usability specifications and measures. 1993.

Biographies



Mike Atyeo
BT Development and
Procurement

Mike Atyeo is currently seconded from the Human Factors Unit at BT Laboratories to provide consultancy to BT Marketing. He has an honours degree in Psychology, and a Masters degree in Computer Science. He joined BT in 1985, having previously worked at NCR (now AT&T Global Information Systems) and Nixdorf on electronic point-of-sale and electronic funds transfer systems.



Rob Green
BT Personal
Communications

Rob Green joined BT in 1987, working on local network modernisation within Marketing. He has since worked on most aspects of marketing within BT, including advertising, promotion, and strategy. He moved to Marketing Services Division in 1990, looking at national advertising expenditure and joined the Portfolio Development team in 1993.

David Cotter, Ian Hawker, Alan Hill and Vivek Tandon

New Network Infrastructures for the 21st Century

Part 2—Ultra-fast self-routing networks

Future multimedia applications will demand radically different networks from those of today. This concluding part of a two-part article considering futuristic network designs describes an integrated network based on optical-fibre transmission of short fixed-length packets, with optical processing and switching at the network nodes.

Meeting the Customers' Needs in the Age of Multimedia

Part 1 of this article¹ described networks of broadband circuits between customers established and controlled by the allocation of optical wavelengths. This allows customers access to a broadband photonic network infrastructure that can support future advanced services. However, the underlying process to set up and close down calls, based on network signalling, has its roots in narrowband telephony. Further into the future it is anticipated that an entirely new approach will be needed. Multimedia computer applications will evolve to provide a slick and seamless interface to the network. These applications will provide instant 'point-and-click' ability to interact with distant users and machines, to browse and interrogate distant databases, to send and receive high-quality images, sound and moving pictures with negligible delay. There will also be increasing computer-to-computer interaction with little human intervention (for example, *software agents*). The network of the future will come to resemble, in function, a large number of networked computers, rather than the historical telecommunications picture of the network as an assembly of interconnected switching and transmission equipment controlled by signalling systems. These new applications and services will generate bursty telecommunications traffic having the following characteristics and requirements:

- short message durations (small fractions of a second);
- long interaction periods with human immersion in real and computer-generated environments;
- network latency dominated by speed-of-light transmission delay;
- extremely rapid call set-up and close-down;
- wide-ranging and rapidly-fluctuating bandwidth requirements;
- wide-ranging message destinations;
- low predictability;
- high location mobility.

It is likely that a future network capable of supporting such extreme traffic diversity, with sufficiently fast reaction times and high reliability, will be radically different from the networks of today and will:

- be composed of primitive elements with simple functionality;
- have simple control systems that are minimal, localised and autonomous;
- exhibit a degree of self-organisation;
- be based on a pragmatic combination of photonics and high-speed electronics;

In this new type of network there will be many more 'nodes' compared to a conventional network, but the nodes themselves will be primitive with very low functional complexity

- have little or no software base;
- have higher intelligence distributed to the periphery.

A radically new network architecture would obviously be required to satisfy these stringent requirements in the future. Based on current knowledge of non-linear optical physics, fibre transmission, photonic devices, and traffic models the following prognosis has therefore been developed.

Network Outline

The network will be capable of spanning a large regional, national or international area, operating as a high-level core interconnection. It will be based on optical-fibre transmission of short fixed-length packets or cells, and photonic processing and switching will be used at the network nodes. The cells will consist of a burst of ultra-short (~1 ps duration) optical pulses at instantaneous rates as high as 100 Gbit/s. The cells will travel across the network from source to destination without translation or alteration, apart from optical regeneration when needed to overcome the impairments and limitations of transmission. Destination address information, coded within an optical header on each cell, will allow the cell to find its own way through the network, exactly in the manner of a letter carried by the postal service (so-called *connectionless* transport). Yet the network will be capable of handling delay-sensitive 'connection-oriented' traffic, such as voice.

In this new type of network there will be many more 'nodes'—positions where traffic steering or processing takes place—compared to a conventional network, but the nodes themselves will be primitive with very low functional complexity. A key concept in the network design is that each of these primitive nodes will be autonomous, will use simple local controls, and will be capable of operating independently on each cell in sequence, one at a time.

These features will allow the network as a whole to respond very rapidly to traffic demands and fluctuations. In fact, for each cell that arrives, the primitive node will perform just one simple binary steering action in answer to some question, such as: is this cell addressed to that destination or not? should this cell go this way or that? should this cell be merged into that traffic stream or not? Some of the decisions will be preprogrammed, deterministic responses based on the header information alone. In other cases, the node will make decisions taking into account instantaneous local traffic loadings, using simple algorithms. It is this feature which will allow the network nodes to exhibit some degree of 'collective intelligence'. The network will have the ability to self-organise—optimise routings, distribute traffic loadings, provide alternative routings; in effect, to respond in real time to the myriad of changes and fluctuations that occur. The tasks of switching and contention resolution will become a communal, collective activity given over to the whole network.

Central network control tasks will be reduced to organising responses to occasional events on a much slower time-scale than traffic fluctuations; for example, a reported cable break or equipment failure will result in instructions to some network nodes to close off certain paths and redistribute traffic loadings.

Building Blocks

The network will use ultra-high-speed transmission in monomode fibre with optical amplifiers and all-optical resynchronisation and regeneration as appropriate. In major research laboratories, the bit rates that can be achieved using optical time-division multiplexing are being pushed ever higher. For example, recent experiments at BT Laboratories² have demonstrated 40 Gbit/s transmission over distances as great as 205 km; Figure 1 shows an eye diagram of the received 40 Gbit/s data. Researchers at NTT Laboratories³ have recently achieved 100 Gbit/s transmission over 200 km.

Yet further transmission distances should become possible in future by using all-optical pulse regeneration and 'shepherding' techniques to reduce timing jitter⁴. There is scope for using a modest degree of wavelength-division multiplexing to increase further the transmission capacity of the fibre links. In that case, passive multiplexers and demultiplexers would be used to separate the various wavelength channels at the network nodes to allow the channels to be processed independently.

The other key photonic components needed to realise the network are:

- lasers able to generate high-quality pulses of duration ~1–2 ps;

Figure 1—Eye diagram of the received signal in a 40 Gbit/s transmission experiment




Figure 2—Oscilloscope traces for semiconductor optical AND gate

- 2×2 optical routing switches capable of reconfiguration in times of ≤ 1 ns;
- planar optical delay-line circuitry with precision better than 1 ps;
- ultra-fast optical AND gates with resolution ~ 1 ps.

All of these components are at the stage of development in research laboratories, and prototypes have been successfully demonstrated. Figure 2 shows a result from a recent demonstration⁵ of a semiconductor optical AND gate device that is capable of operation at speeds greater than 100 Gbit/s. Compared to digital electronics, photonic logic devices are much less developed; they tend to be larger, poorly integrated and more expensive. Nevertheless, they have an important role to play since they are capable of operating at bit rates far exceeding the physical limitations of electronics.

The general approach that has been adopted in designing the signal-processing nodes for an ultra-fast network is to use a small number of photonic devices with simple functionality, such as the optical AND gate, for any operations that must essentially be performed at the picosecond bit level, leaving digital electronics to handle the other cell-level processing functions that can be performed at lower speed (0.1–1 GHz clock speed). Thus, for example, photonic processing with picosecond speed of response would be used for bit-level recognition of the cell header address, while electronic hardware would be used to make routing decisions and control the opto-electronic switches.

Ultra-Fast Cells and Network Timing

At the periphery of the self-routing network, incoming data carried on conventional (lower bit rate) data channels will be captured during the process of creating fixed-length ultra-fast cells. This will use the techniques




Figure 3—Time-slot diagram showing time windows for a fixed-length ultra-fast cell

of optical pulse replication, modulation and bit-interleaving developed already for multi-gigabit transmission. Cells leaving the network will be 'decompressed' using photonic switching to convert the data back to lower-rate channels.

The network will use a universal clock to provide network synchronisation at the cell level. Figure 3 shows an example of the partitioning within a time-slot. Here the optical cell consists of a burst of ultra-fast optical pulses representing about 440 bits (420 bits for the payload and application header, plus around 10–20 additional header bits to allow self-routing of the cell in the network) at an instantaneous rate of 100 Gbit/s. The time-slot also contains a switch band of ~ 1 ns, allowing time for the reconfiguration of routing switches, and time guard bands. In this example, the network clock has been chosen to be the standard synchronous digital hierarchy (SDH) STM-1 rate of 155.52 MHz (period 6.43 ns). The

time-averaged data rate for the cell stream is therefore

$$420 \times p \times 155.52 \text{ Mbit/s} = p \times 65.3 \text{ Gbit/s},$$

where p is the traffic load (probability of occupation of a time-slot). The network clock provides synchronisation at the cell level only, not at the picosecond bit level. It will be impractical to maintain synchronisation of cells with picosecond precision throughout the network. Instead, some simple techniques have been developed that will allow timing information to be extracted from the individual cells, on a cell-by-cell basis, wherever required.

Binary Steering and Routing Strategies

As described earlier, the network will contain a large number of primitive nodes whose purpose is to perform some simple binary steering operation on cells at the junctions of the

Figure 4—Layout of a binary cell routing node

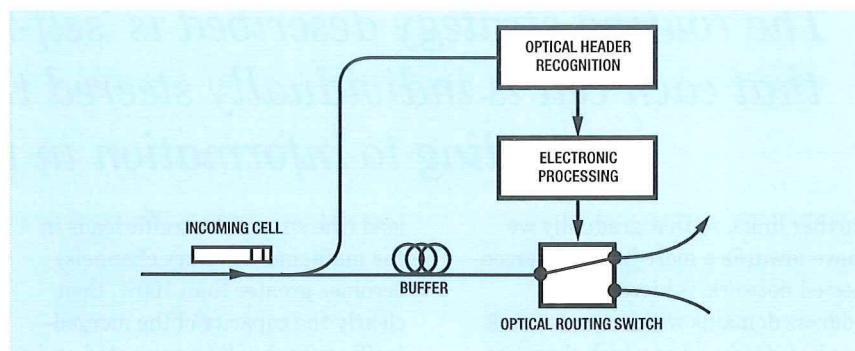
network. The most fundamental operations are route selection (making a choice between two available routes) and merging (a junction that combines two traffic streams). The simplest types of node will perform these steering operations in a preprogrammed deterministic fashion. An example of this type is a node which will direct a cell to either one of two output paths according to the destination address (or a portion of it) carried in the cell header field (Figure 4).

The optical header recognition unit must read and recognise binary words encoded as a short ultra-high-speed train of picosecond optical pulses, and a simple way of doing this has been developed that requires only a single optical AND gate⁶.

The second fundamental operation that is required is the ability to merge two independent cell streams (a main stream and a tributary) into a single flow of traffic, as depicted in Figure 5. A photonic switch fabric has been developed that is capable of performing this function with low cell-loss probability (provided the input traffic load is not excessive), requiring only a small number of optical components and modest electronic hardware⁷.

By having the ability to perform these two fundamental operations—binary route selection and merging—it is possible to construct the simplest kind of network: a unidirectional ring, as shown in Figure 6. This is clearly a rather limited type of network, however, because the average traffic capacity available at an access point is just the transmission capacity of the ring divided by the number of access points.

A network with complexity one step higher than the unidirectional ring would be a dual bus or dual ring, with two counter-propagating traffic streams. The nodes within such a network must decide the most appropriate direction for a cell to take, based on the destination address in the cell header. A binary routing algorithm for such a self-



routing network can be constructed based on the recognition of certain short words or fragments of words found in the address field. The recognition process can be performed at the necessary ultra-fast speed using simple optical-AND techniques. For example, in a dual-bus network architecture, there might be 20 destinations. That number of distinguishable addresses requires 6 bit address words if restricted code word sets required for the optical-AND recognition technique⁶ are used. The trick then is to work out the minimum number of optical-AND operations that must be carried out on fragments of those 6 bit words to determine the correct direction of onward propagation of a cell, assuming the destination addresses are arranged in an appropriate numerical sequence along the bus. It turns out that, in this example of 20 destination addresses each six bits long, only one or two (or very exceptionally three)





Figure 5—A binary node that allows two input cell streams to merge

optical-AND operations need to be performed before the correct propagation direction is determined.

Suppose that we have a ring architecture and that a node in the ring can perform a simple recognition operation on the header of a cell, the result of which indicates that the cell is addressed to somewhere on the far side of the ring. Then, rather than send the cell off on its way laboriously step-by-step around the circumference of the ring, it would be more effective to steer the cell directly onto a link connecting opposite sides of the ring, thus bringing the cell more quickly towards its destination. This idea can be extended to introduce

Figure 6—An elementary ring network for ultra-fast cell transport



The routing strategy described is 'self-routing' in the sense that each cell is individually steered through the network according to information in the header

further links, so that gradually we move towards a more highly interconnected network. A hierarchy of address domains within the network can be introduced in which there are a relatively small number of distinct addresses within each domain. The address of a domain in each successive level in the hierarchy would be represented by a short portion of the header address field. This is quite like the different lines of the address on a letter in the postal system: the name of the county, town, street and so on. The key point here is that the address is carried intact by the cell as it travels throughout the network, and the task of reading a portion of the address (using ultra-fast optical word recognition) and making an appropriate binary routing decision is made as simple as possible.

Self-Organisation

The routing strategy described is 'self-routing' in the sense that each cell is individually steered through the network according to information in the header. This procedure is a deterministic one, preprogrammed by the network designer. However, as mentioned earlier, traffic patterns will be increasingly erratic and unpredictable in the future and there may be no economically-viable network configuration that is appropriate to handle all traffic demands. A solution to this problem can be found by allowing the network itself continually to select traffic routings according to the instantaneous traffic demands upon it; in other words, to exhibit a degree of self-organisation. If this self-organisation is to occur on the necessary time-scale of extremely rapid fluctuations in traffic loads, it will be essential for the decision making to be performed locally and autonomously.

An elementary degree of self-organisation can be obtained using the merging junction shown in Figure 5. As described earlier, if the input traffic load is not excessive, very low cell-loss probabilities can be obtained. If, however, the total input

load (the sum of the traffic loads in the main and tributary channels) becomes greater than 100%, then clearly the capacity of the merged-traffic output will be saturated and a substantial amount of traffic will begin to emerge from the overflow. Instead of merely allowing the overflow traffic to be lost, it can be carried away by a separate optical path, a strategy known as *deflection routing*. The deflection route may lead towards the same onward node as the main route, but follow a different geographical path, or incorporate some substantial length of extra fibre, to introduce a relative propagation time delay. If this time delay is longer than a typical burst of traffic congestion, the overflow traffic can be successfully merged with the main traffic at the distant location. Alternatively the deflection path might lead to another part of the network. In this way, a network of alternate routings can be built across the network to overcome local congestion.

This idea of using a 2×2 junction as a rudimentary traffic controller can be developed further by introducing a modest amount of additional intelligence. Cell detectors could be located in each of the two input (or output) paths at a traffic-merging node to provide fast real-time monitoring of the local traffic loads. This information could be used to adjust the operation of the node. Even very elementary algorithms can provide powerful functionality, and could be implemented using simple electronic hardware (a small number of fast logic gates). For example, an intelligent node could be programmed to deflect traffic once the input load exceeds a preset level. This feature could be valuable, for example, in ensuring sufficient spare capacity on the outward link so that other traffic-merging operations further downstream can operate effectively. Another possibility would be for a node to provide a traffic-balancing function, in which the traffic on two input lines would be

divided between the two output lines in a chosen ratio.

In general the approach will be to balance out traffic loadings across the network as far as possible and avoid operating individual links at too high a capacity, to avoid congestion 'hot spots' and maintain sufficient degrees of freedom for effective network self-organisation. This is a practical example of the key concept of trading off some of the available transmission bandwidth in return for increased simplicity of network control and management.

Networks

Figure 7 shows a national network in which a self-routing 'core' is used to interconnect a large number of circuit-switched local access networks. An individual customer who wishes to use the self-routed network for some application will be connected via the local access network to a feeder (a place where traffic can be directed into and out of the network). At the feeder, traffic originating from the customer in some standard format (for example, asynchronous transfer mode (ATM) cells) would be converted into ultra-fast optical cells for self-routed transmission across the network. Each cell header would include the address of the feeder point at which that cell is to exit the network, for onward relay to the final destination. At the exit feeder, the ultra-fast cells would be decompressed and converted back to the standard format. As far as the customer is concerned, the network is a transparent link. For example, an ATM application would specify the virtual paths and circuits as necessary for the links into and out of the feeders, but everything between the feeders, within the ultra-fast self-routed network, would be invisible to the customer. In fact, it seems likely that new applications will evolve that will make best advantage of the ability of this network to establish instant high-bandwidth virtual connections.

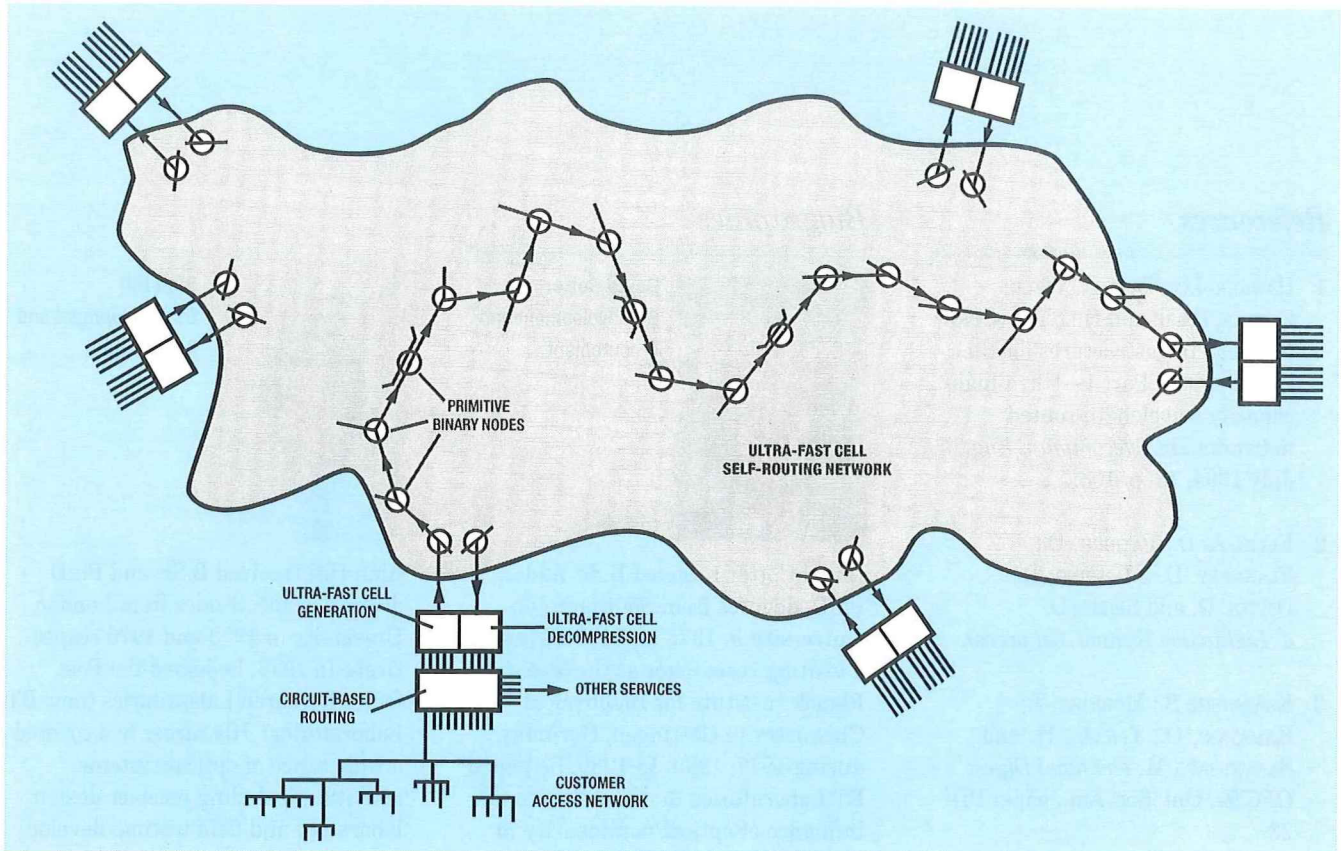


Figure 7—Layout of a network that provides fast broadband connections between customer-access networks via an ultra-fast self-routing core

A preliminary study indicates that the task of reassembling messages from the incoming stream of cells (possibly in disordered sequence) in real time is within the capability of a modest processor in the customer's equipment.

Access to the ultra-fast network would be limited only by the capacity of the feeders. The feeder capacity would be strictly dimensioned to avoid overall network congestion. The core network capacity is given simply by the product of the feeder capacity and the number of feeders. For example, 3000 feeders (comprising an average of 60 feeders at each of 50 major population centres), each operating at ~30 Gbit/s payload capacity (100 Gbit/s peak line rate and 50% traffic load), would provide a core capacity of ~100 Tbit/s. At any instant this is sufficient capacity for connectionless service to one million customers at an average of 100 Mbit/s (or 10 million customers with 10% instantaneous network utilisation per customer). In this scenario each of the population centres would be served by 60 optical fibres, or fewer fibres if a modest amount of wavelength-division multiplexing is used.

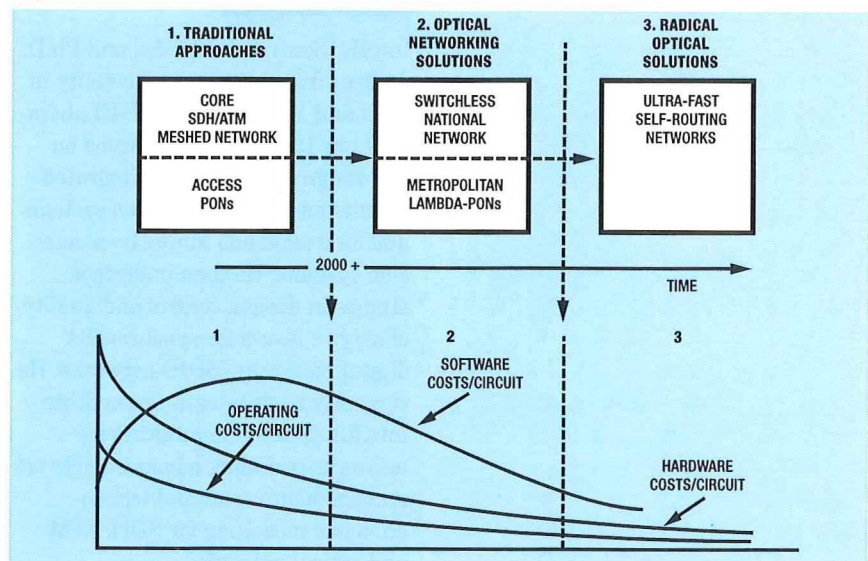
Summary

The futuristic developments outlined in the two parts of this article are driven by the need to reduce hardware, software and operating costs for telecommunications networks (Figure 8). Approaches have been described that are very different from the traditional view of a meshed SDH/ATM core network linked to separate PON access networks. These

are replaced by integrated solutions using optical network technology to provide bandwidth transparency, simplified switching (either simplified/distributed or eliminated) and simpler network management.

Research into these highly futuristic network designs is ongoing. These new network concepts differ radically from today's approaches, and yet are clearly possible in the early part of the next millennium.

Figure 8—Network evolution trends



References

- 1 HAWKER, IAN; TANDON, VIVEK; COTTER, DAVID and HILL, ALAN. New Network Infrastructures for the 21st Century. Part 1—Ultra-high-capacity wavelength-routed networks. *Br. Telecommun. Eng.*, July 1994, **13**, p. 103.
- 2 ELLIS, A. D.; PATRICK, D.; FLANNERY, D.; MANNING, R.; DAVIES, D. and SPIRIT, D. *J. Lightwave Technol.* (in press).
- 3 KAWANISHI, S.; MORIOKA, T.; KAMATANI, O.; TAKARA, H. and SARAWUTARI, M. *Technical Digest OFC'94*, Opt. Soc. Am., paper PD-23.
- 4 WIDDOWSON, T.; MALYON, D.; ELLIS, A.; SMITH, K. and BLOW, K. *Electron. Lett.*, 1994, **30**, 990.
- 5 NESSET, D.; TATHAM, M.; WESTBROOK, L. and COTTER, D. *Technical Digest Photonics in Switching*, 1994 (in press).
- 6 COTTER, D., and COTTER, S. C., *Electron. Lett.*, 1993, **29**, 945.
- 7 COTTER, D., HUNTER, D. and ANDONOVIC, I. (submitted to *Electron. Lett.*)

Biographies



David Cotter
BT Development and
Procurement

David Cotter received B.Sc. and Ph.D. degrees from Southampton University in 1973 and 1976. He was a visiting researcher at the Max-Planck Institute for Biophysical Chemistry in Göttingen, Germany, during 1977–1980. In 1980, he joined BT Laboratories to research into the influence of optical nonlinearity in telecommunications systems and ultrafast optical nonlinearity in semiconductors, organic materials and glasses. His current interests are the development of radically new network solutions for the future, using ultrafast photonic technologies. In 1993, he was elected a Fellow of the Institute of Physics.



Ian Hawker
BT Development and
Procurement

Ian Hawker received B.Sc. and Ph.D. degrees from Leicester University in 1973 and 1976. He joined BT Laboratories in 1976 where he worked on the design of high-speed integrated circuits for TAT-8 submarine systems and for inland 565 Mbit/s transmission systems. He then undertook studies in design, control and quality-of-service issues for synchronous digital hierarchy (SDH) networks. He currently leads a team researching into future telecommunications networks including advanced network management issues and techno-economic modelling for SDH, ATM and optical networks.



Alan Hill
BT Development and
Procurement

Alan Hill received B.Sc. and Ph.D. degrees in Electronics from London University in 1973 and 1976 respectively. In 1973, he joined the Post Office Research Laboratories (now BT Laboratories). His career has covered a wide range of optical systems activities, including receiver design, laboratory and field testing, development of micro-optic components such as connectors, switches, tunable transmitters/receivers and wavelength-division multiplexers, studies of linear and non-linear optical crosstalk and the development of advanced optical switching and distribution networks. He has more recently been involved in the broadband upgrade of telephony passive optical networks (TPON/BPON) by WDM, including responsibility for BT Laboratories' activities in the RACE II MUNDI project for fibre-to-the home and building.



Vivek Tandon
BT Development and
Procurement

Vivek Tandon received his B.Sc. degree in Physics at Imperial College (University of London). He joined BT in June 1993 and is researching the development of future revolutionary packet and non-packet based optical networks. He is also pursuing a part-time Ph.D. at the University College of London developing a broadband wavelength-division-multiplexed passive optical network (Lambda-PON).

Professor Alistair G. J. MacFarlane

The Future of Learning

Technology has an indispensable role to play in creating the learning support environments of the future. To design an effective learning support system requires an appropriate model of the learning process. It is believed that the principles that underly machine-aided support could be adapted for use in a variety of fields of learning.

Introduction

All learning involves a dialogue between imagination and experience. Learning of any sort is an active, dynamic process, and can only take place effectively given a suitably supportive environment and suitable media for expression. These media should provide both a rich resource and a flexible stimulus. One of the traditional media—the written document—has been transformed over the past decade by the technology of word processing. Slowly but surely, all documents in most enterprises are now being produced by using increasingly sophisticated word-processing systems.

The importance of multimedia techniques, and the opportunities which are offered for all educational and creative work, arise simply from the fact that it is now possible to manipulate and store virtually all forms of media—audio, video, text, pictures and whatever. With the same flexibility, information can be assembled by cutting and pasting video, audio and text into animated documents, films or whatever. Together with these media processing systems has come an astonishing development of recording and storing devices. Video recorders are now a relatively inexpensive tool; cameras can record directly into digital form for computer storage and hence subsequent editing; and compression techniques and compact-disc storage can handle the large amounts of information required for storing a variety of media.

The media revolution is underpinned by developments in storage and processing technology. The speed of processing and the storage required for combining large amounts of media into presentations will both soon be cheaply available. Multimedia processing systems will then spread in the way that word process-

ing has spread. This will happen over about the same time as the displacement of the typewriter, if not in quite the same pervasive way. Because it is something quite new, the implications are enormous. What is being offered will be, for most of us, an entirely new range of media for creation and expression.

In all countries, which are determined to maintain or sustain reasonable levels of economic growth and improve their quality of life, there will be a steady increase in the use of technology in education and training. There are several reasons behind this assertion.

- The inexorably increasing demand for more and better-quality education and training at all levels will require the creation of highly-effective and cost-efficient environments for the support of learning.
- The steady development of information technology and the spread of consumer electronics and communications networking will make the use of high-grade information technology (IT) inevitable for the creation and support of learning environments.
- The opportunities for innovation which the availability of such environments will create, and the associated necessity to take careful account of the cognitive aspects of their design, will rapidly increase both the need and the opportunity for high-quality research to underpin the design of such systems.
- The proliferation of the technology of self-paced tutor-supported distance learning will both revolutionise the use of training in industry and the provision of professional training throughout a

By using technology, powerful environments could be created for the flexible and efficient support of the learning process

working life, and so will open up significant commercial opportunities both for the provision of courseware and for the delivery platforms on which it will run.

Thus there are important implications for all forms of continuing professional education and training, and for the flexible and adaptable provision of facilities for learning throughout life. By using technology, powerful environments could be created for the flexible and efficient support of the learning process.

The key concepts underlying the creation of such environments, together with a specific example, are now outlined to illustrate the potential of technology for supporting learning in a flexible and adaptable way. In order to consider how IT can best be deployed to support learning, the processes of learning and teaching are first examined in some detail.

The Structure of Learning Support

Learning is an interactive and dynamic process with imagination driving action in exploring and interacting with an external environment. *Concepts* are abstracted from experience and result in comprehension. Effective action is achieved via *schemata*—the rules and procedures which are assembled to guide effective action. Knowledge is built up via a cyclic process of guided exploration, by informal trial and error. Learning results from the progressive development and refinement of concepts and schemata. Learning leads to the acquisition of coherent frameworks of reasonable beliefs together with the necessary skills to put them to effective use.

Knowledge comes in *chunks*. Our limited raw-information-processing capability requires us to interact with only a severely limited part of our environment at any specific time. *A fortiori* when we learn we can only learn about, and acquire a facility in

dealing with, a highly limited version of part of the world. The severely limited, circumscribed and simplified situations which we can attend to, or learn about, at any specific time is called a *microworld*. This idea is central to the treatment of machine-supported learning which is given here. As we learn by interacting with our environment, day after day, year after year, we weave, from our experience, an intricate web of beliefs, concepts, descriptions, prescriptions, rules and procedures, facts and dodges. We weave microworlds into macroworlds, constantly modifying, stitching, repairing and creating our abilities to understand and to cope. This process is active, constant and never-ending. The metaphor of weaving a web of belief has been discussed in detail by Quine and Ullian¹ and is of great value in describing the learning process.

In interacting with the world, we encounter and use information in two ways². One way characterises *the world as sensed*—it provides, for example, the information needed for identifying and characterising objects and for building models of them. The other way characterises *the world as acted upon*—it provides the information needed for an action to have a desired effect or, for example, to produce or generate objects having a desired set of characteristics. Thus, information can be both *descriptive* and *prescriptive*—one use of information is associated with specifying *perception*, the other is with specifying *action*.

Descriptions and prescriptions are the warp and weft of our web of understanding. In interacting with a microworld, we need both to comprehend it and to be able to act effectively upon it. To comprehend it, we need a necessary set of concepts, and to act upon it we need a necessary set of schemata. To understand a microworld is to understand the set of concepts associated with it and to be able to use the set of schemata defining how to act upon it.

Concepts

Skemp³ studied the psychology of learning mathematics which is very useful in considering learning-support mechanisms. His approach is built on the following postulates for concepts.

- A concept requires for its formation a number of experiences which have something in common.
- A concept is a way of processing data which enables the user to bring past experience to bear on a current situation.
- Concepts of a higher order than those which learners already have cannot be communicated to them by a definition, but only by arranging for them to encounter a suitable number of examples.
- If any of these examples are, or involve, concepts, it must be ensured that these are already formed in the mind of the learner.

Schemata

Skemp's approach also makes use of schemata, which he takes to be the general psychological term for mental structures or mental models. According to Skemp, a schema has two functions in the learning process:

- it integrates existing knowledge; and
- it is a mental tool for the acquisition of new knowledge.

His approach to learning mathematics is based on the proposition that to understand something means to assimilate it into an appropriate schema. Since understanding is a subjective experience, this places a pragmatic emphasis on the unique role of the individual. For a characterisation of knowledge and learning which can be used in the development



The conceptualisation of a mathematical process

Learning and its Phases

The following point of view is taken in approaching the analysis of systems for the support of the learning process.

- Learning is an active feedback process in which concepts are acquired, associated with appropriate schemata and tested in action.
- Learning is a pragmatic process, centred on the individual.
- Learning requires interaction with objective knowledge structures, and results in the generation of new subjective schemata.
- The effectiveness of learning can only be assessed in terms of its practical consequences.

The development of thorough conceptual understanding involves a series of learning *phases*—preparing to tackle the relevant material, acquiring the necessary information, relating it to previous knowledge, transforming it through establishing organisational frameworks within which to interpret it, and so developing personal understanding. If this process is to work effectively, teaching—however it is delivered—must be *designed* to support these phases of learning. The support required can be described in terms of necessary *teaching functions* which to some extent parallel, but also overlap, the phases of learning. These functions include:

- orientating—setting the scene and explaining what is required;
- motivating—evoking and sustaining interest;
- presenting—introducing new knowledge within a clear, supportive structure;
- clarifying—explaining with examples and providing remedial support;

of machine-supported learning systems, one can take a somewhat narrower view of schemata, and adopt the following postulates, in terms of a narrow and highly specific domain of experience; that is, of a microworld:

- a schema is a coherent, integrated, logically connected framework of descriptions and prescriptions for a specific microworld;
- schemata encapsulate understanding of microworlds in terms of giving prescriptions for effective actions in that microworld;
- a schema provides a user hand-book for that microworld;
- schemata link concept and action, knowing and doing, observation and explanation for specific microworlds;
- assimilation of schemata is manifested by an ability to give explanations of a microworld and to solve problems posed in terms of it; and
- schemata are acquired by interaction with a microworld and the

objective knowledge characterising that microworld, and its organisation into prescriptions for purposive, effective action with reference to that specific microworld.

A microworld is a greatly simplified but logically coherent system which represents a finite and specific part of the world, and which is created for the purposes of teaching or training.

Understanding

Understanding of a microworld is manifested by an ability to:

- use the concepts assimilated and the schemata created in the learning process to interpret data successfully, explain sets of related events and solve problems posed in terms of the microworld;
- cope with new situations described in the microworld;
- explain new situations arising in the microworld; and
- act in these new situations with satisfactory consequences.

Figure 1—Roles of teacher and learner

- elaborating—introducing additional material to develop more detailed knowledge;
- consolidating—providing opportunities to develop and test personal understanding; and
- confirming—ensuring the adequacy of the knowledge and understanding reached.

After the initial phases of preparing for learning, the conventional process of acquiring new knowledge begins with the teacher *presenting* appropriate course material. The new knowledge has to be carefully selected to ensure relevance and potential interest, then presented in a way that helps the student both to relate it to prior knowledge and to see a clear logical structure within it, as a first step towards establishing a personal organising framework. Then the process of *clarifying* begins. Through explanations and examples, students are encouraged to begin developing their own personal understanding of the topic. If the prior knowledge is inadequate, remedial support will be required at this stage, to allow a firm base on which to build explanations and further clarification. Once the initial grasp of the material is sufficiently firm, opportunities have to be provided for *elaborating* the knowledge by examining nuances of meaning and by incorporating more detail and additional examples or evidence. Thereafter, the knowledge requires *consolidating*, by encouraging its application to new contexts, and periodically reviewing what has already been presented. Ultimately, there needs to be a final consolidation that allows the students to integrate the course as a whole; and that is often linked, through assessment procedures, to the final teaching function—confirming that knowledge and understanding have reached an appropriate level. Such confirmation is involved in certifying standards to the outside world, and

as part of the quality controls to ensure that the teaching has been effective.

Within the conventional teaching methodology, the initial stages of orientating, motivating, presenting and explaining or clarifying are carried out through lectures. Further clarification and remedial support is obtained from textbooks or through tutorials which also involve elaborating and consolidating. Laboratory classes and field work introduce additional knowledge and skills, together with opportunities for consolidation and elaboration. The additional reading suggested by lecturers and tutors continues the process, while much of the consolidation comes from worked examples and coursework essays. In all areas of study, assessment requirements are used to encourage consolidation through periodic reviews and the thorough revision process which precedes examinations. Students' progress needs confirmation through comments on coursework and the results of periodic tests, while degree examinations are used to confirm formally the levels of skill and understanding reached by students. Finally, course evaluations are required to confirm the quality of the teaching.

The provision of teaching and the support of learning should be seen in terms of a complex interacting system. Assessment is only one of the factors influencing learning. The outcomes of learning depend on the combined effects of the whole *learning environment* (teaching, discussion classes, resource materials, and assessment

procedures) provided by the institution and its courses. The provision of an effective and economical system requires a careful analysis of requirements and functions both at institutional and course levels.

While the conventional teaching methods continue to fulfil these functions, it is now necessary to consider whether there should be widespread adoption of more efficient and cost-effective ways of encouraging student learning. There needs to be a re-examination of both the purposes and the techniques involved in conventional teaching methods, and more extensive use made of new methods that support the additional transferable skills now being required. In deciding how to provide the most efficient and effective form of teaching, it will also be useful to keep in mind three general questions:

- How can the knowledge be structured and presented most effectively?
- What forms of interaction are most appropriate to develop understanding?
- In what ways can this interaction best be managed?

This approach to teaching and learning is illustrated by Figure 1. In it the teacher no longer functions as a transmitter of information but allows the learning support functions to be supplied by a machine, as illustrated by Figure 2. An important implication of this form of delivery system struc-

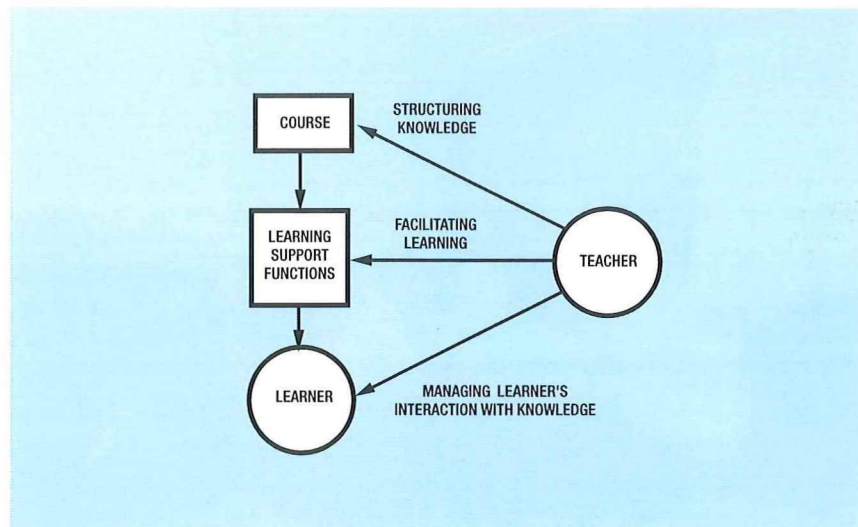
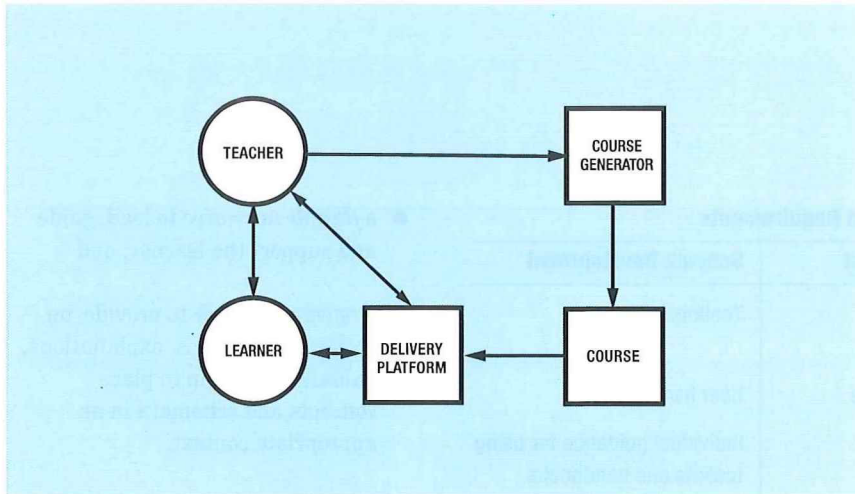


Figure 2—Machine-supported learning



ture is that *remote* tutoring can be provided, as illustrated by Figure 3.

Intensely-Supportive Learning Environments (ISLEs)

The metaphor of an environment is very useful when discussing support for the learning process. From this point of view, education is the design, creation and management of environments which support the learning process, and teaching is the management of the learning process within such an environment.

In a narrower sense, one can conceive of creating an intensely-supportive environment for learning by using computer-based systems to aid the human reasoning process, and communication systems which provide flexible access to tutors and to fellow learners. The compound adjective 'intensely-supportive' denotes the

fullest possible exploitation of available computing and information-handling technology. A learning-support environment is a computer-based system that allows interaction with structured, objective knowledge in a way that enables the user to absorb concepts and develop schemata, derive an understanding of their use and to test that understanding by appropriate investigations and exercises. Such an environment should support learning on both an individual and a group basis and allow the users great flexibility in finding their way through the structured, objective knowledge, which is made available to them, yet also allow an instructor to monitor and, if necessary, direct the progress of each individual student. In such a supportive environment, the learner should be able to perform a variety of tasks:

- build representations and models,

- analyse,
- browse,
- search,
- compare and evaluate,
- reason and hypothesise,
- synthesise,
- manipulate and modify,
- experiment,
- catalogue, and
- store and retrieve.

If this wide diversity of tasks and facilities is to be made available in an efficient and flexible way, then the support system must be:

- easily comprehensible to a single individual;
- wide in scope;
- modular, with a manageable number of distinct parts;
- predictable in its behaviour;
- integrated and coherent in the ways in which its different parts relate;
- helpful, with quick and efficient access to relevant information;
- tolerant of errors and supportive in enabling the effects of errors to be easily undone;
- extensible and adaptable; and
- self documenting.

It will be essential in developing such supportive learning environments to achieve an appropriate sharing of tasks between the human and the machine. An intensely-supportive learning environment (which has the convenient acronym, ISLE) must provide:

Figure 3—Remote tutoring

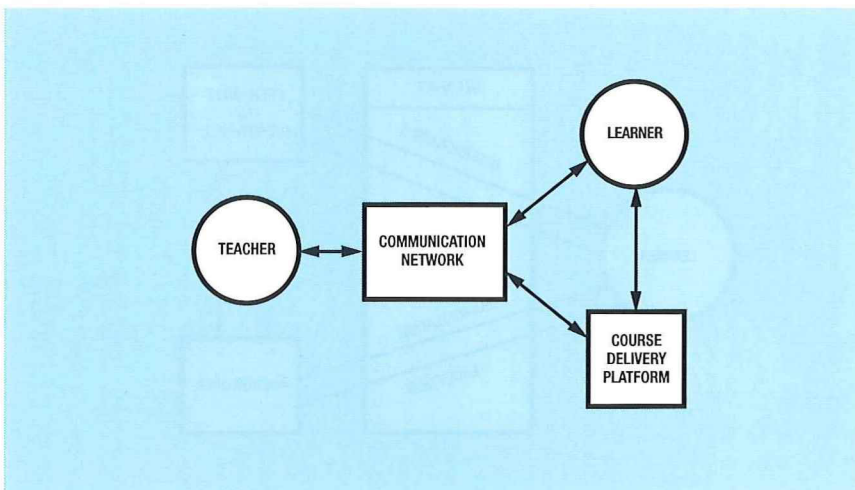


Table 1 Learning Environment Support System Requirements

Data Support	Concept Development	Schema Development
Interactive simulators	Subject knowledge descriptions	Toolkits
Model Builders	Instructional packages	User handbooks
Virtual experiments	Tutorials	Individual guidance for using toolkits and handbooks
Textual material	Dictionaries and thesauruses	Performance monitoring
Video and audio material	Synoptic overviews	Tests and evaluations
Databases	Plans and strategies for learning	Prompts and advice
	Tests and evaluations	
	Prompts and advice	

- a *flexible narrative* to lead, guide and support the learner; and
- a *reference system* to provide, on request, definitions, explanations, guidance and help to place concepts and schemata in an appropriate context.

The way in which the various parts of a learning support system fit together is shown in Figure 6, with appropriate names for the key system components. An *instructor* exercises overall control in a feedback loop driven by the *evaluation* of understanding by the *learner*. The instructor works through *narration* and *demonstration*. The narration guides concep-

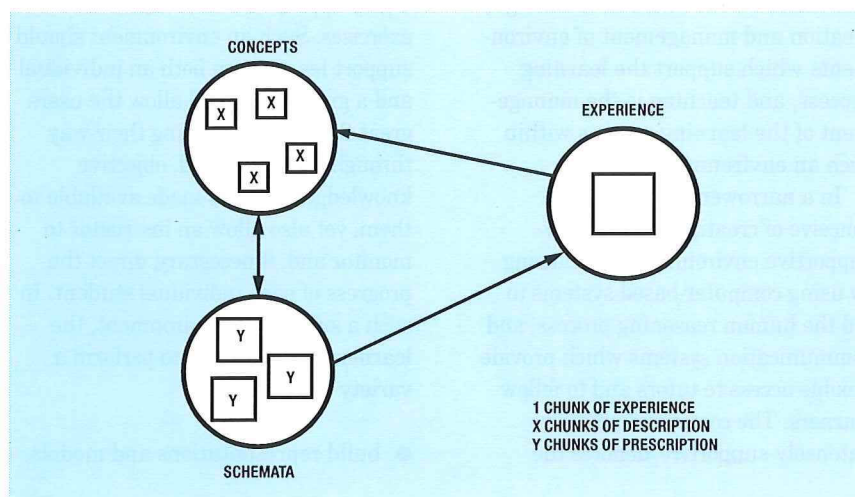
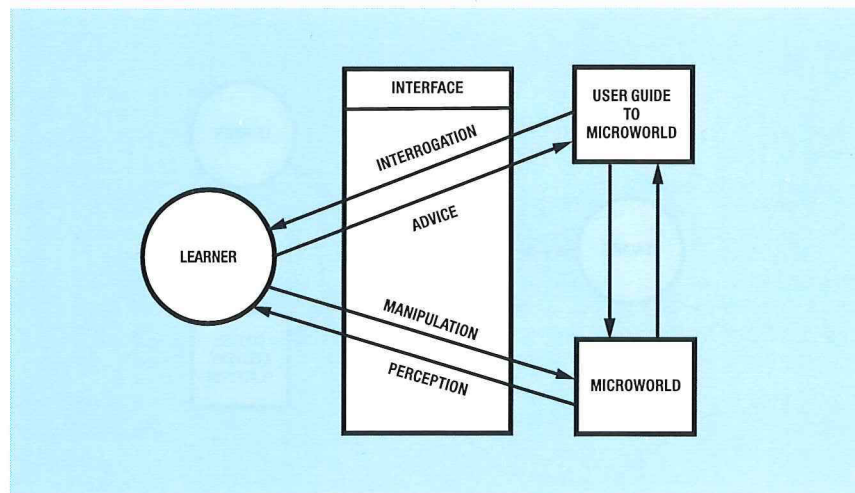
- a data-support system providing powerful data-handling facilities;
- a concept-development support system giving flexible access to structured, objective knowledge; and
- a schema-development support system underpinning the students' generation of new knowledge schemata, and guiding their testing in a variety of uses.

Table 1 shows the individual requirement that each of these support systems would provide.

Associated with the formalised chunks of experience (microworlds) are chunks of concepts and chunks of schemata; as indicated by Figure 4, the numbers of these chunks will normally be different.

Figure 5 illustrates the learner's interaction with a microworld. The overall functional structure of a learning support system can now be sketched out. The overall system must provide, among other things:

- development of an understanding of the *concept set* for the microworld;
- development of an ability to use the *schemata set* for the microworld;

Figure 4—Chunking*Figure 5—Interaction with microworld*

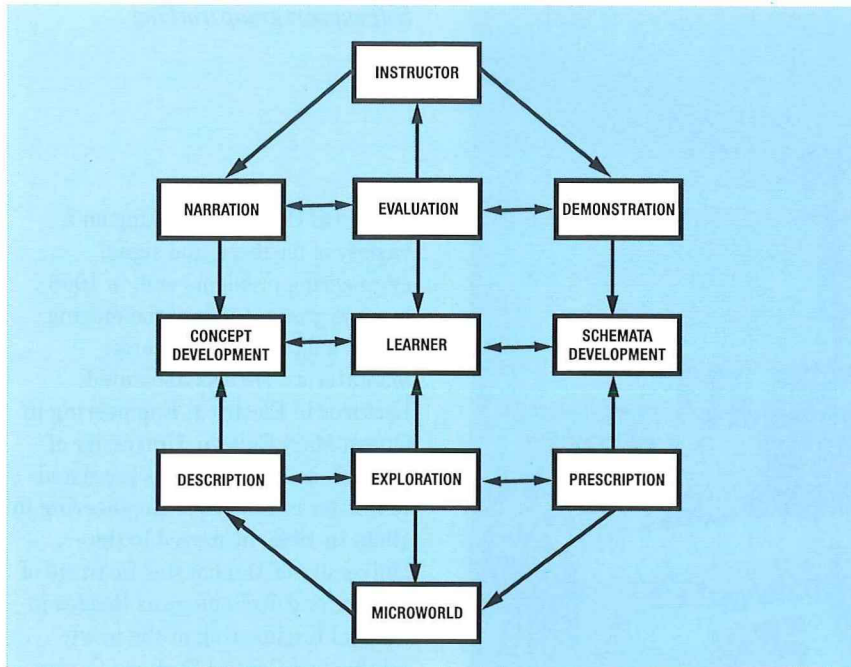


Figure 6—Functional structure of learning support

- requires proficiency in the use of a related set of schemata whose use must first be mastered;
- is expressed in terms of a user guide for effective interaction with a part of the microworld whose properties are being taught; and
- generates the ability to carry out certain tasks whose performance can be tested by evaluation.

It must be emphasised that machine-based learning-support environments are seen as *supplementing*, and not displacing, the essential role of the human teacher. Their use will free the vital human skills for one-on-one and small-group tuition, of which so little is now available at so many levels of education. The use of technology for learning support is likely to evolve in such a way that the group will become the natural learning unit, and pupils and students will increasingly reinforce each other's learning experiences, guided by both their peers and by their teachers.

There are wider implications to what is being proposed:

- powerful learning support environments coupled with remote tutoring could provide small- and medium-sized enterprises with affordable training schemes;
- tutorial support delivered by telecommunications into such supportive environments, based on portable computers in the home or workplace, could be the prime delivery mechanism for advanced professional training; and
- the provision of powerful knowledgeable machines (which is one way of looking at a learning support system), with which one could interact easily, could go a long way towards overcoming the increasing fragmentation of knowledge into highly specialised domains.

tual development and the demonstration guides *schemata development*. The learner *explores* the *microworld*, building up an effective set of *descriptions* and *prescriptions*.

At the beginning of the course of instruction, the narration *orients* and *motivates* the learner, sets the scene, explains what is required and evokes interest by running a set of *scripts*. Linking with the demonstration, the narration also introduces the way in which schemata are developed via *user guides*.

The instructor, driven by the *evaluation* of understanding, controls the *presentation*, *clarification*, *elaboration* and *consolidation* of knowledge. This is done by providing interactive feedback-controlled support for the development of the necessary conceptual understanding and schematic skills for the specific microworld for which the course is developed.

The essence of such a learning support system is that it:

- provides a flexible yet systematic framework for the creation of courses;
- is based on a sound characterisation of the process of learning support;
- can use presentation package methodologies to present concepts;

- can use interactive user-guide methodologies to generate schemata; and
- can use a combination of both these methodologies to knit together an interactive learning support package.

All concepts are part of a linked set of concept frameworks and schemata frameworks. Each concept:

- requires, for its assimilation, a set of prior concepts;
- has an associated set of scripts which develop the necessary understanding and places it in an overall context;
- generates a degree of understanding which can be tested by evaluation; and
- provides the necessary conceptual understanding needed for effective interaction with the part of the microworld whose properties are being taught.

All schemata are part of a linked set of schemata frameworks and concept frameworks. Each schema:

- requires a set of prior concepts to be understood in order that it may be used effectively;

*Self-support group working*

Conclusion

The purpose of this treatment of the key ideas underlying the way in which technology could be used for the support of learning aims to show how, if used in a sensitive and appropriate way, technology could make an invaluable contribution to the development of educational and training provision. A subsidiary purpose is to give an analysis of the process of learning support which, it is believed, has wide applicability.

Large-scale changes to anything as complicated and as important as the country's educational and training systems will be a complex, prolonged and disruptive process. They will have to be carried through in the face of the ever-present realities imposed by costs, space, time, organisational and institutional constraints, and by individual attitudes. But germane to all of this is the key contribution of telecommunications.

Acknowledgement

The material presented here formed part of the 1994 Robbins Lecture, delivered at the University of Stirling on 9 May 1994. The treatment of the

phases of learning draws on work by Professor Noel Entwistle of the University of Edinburgh.

References

- 1 QUINE, W. V., and ULLIAN, J. S. *The Web of Belief*. Random House, New York, 1978.
- 2 SIMON, H. A. *The Sciences of the Artificial*. MIT Press, Cambridge MA., 1982.
- 3 SKEMP, R. R. *The Psychology of Learning Mathematics*. Penguin Books, London, 1971.

Biography



Professor Alistair MacFarlane
Principal and
Vice-Chancellor,
Heriot Watt University

Alistair MacFarlane graduated in electrical engineering from the University of Glasgow, in 1953. He joined the Metropolitan-Vickers

Electrical Company working on a variety of feedback and signal engineering problems and, in 1956, became group leader of the moving target indication and receiver laboratories. He was appointed Lecturer in Electrical Engineering in Queen Mary College, University of London, in 1958, and was promoted to Reader in Electrical Engineering in 1965. In 1966, he moved to the University of Manchester Institute of Science and Technology as Reader in Control Engineering at the newly established Control Systems Centre, becoming Professor of Control Engineering in 1969. He was elected to a Chair of Engineering in the University of Cambridge in 1974, and became Head of the Information Engineering Division of the Engineering Department at Cambridge. He took up his present position as Principal and Vice-Chancellor of Heriot-Watt University, Edinburgh, in 1989.

Professor MacFarlane is an Honorary Fellow of Selwyn College, Cambridge, a Fellow of the IEEE, of the Institute of Measurement and Control, and of the IEE; and Consultant Editor of the *International Journal of Control*. He was elected to the Royal Academy of Engineering in 1981, a fellow of the Royal Society in 1984, a Fellow of the Royal Society of Scotland, and was awarded the CBE in 1987. The IEE awarded him its Achievement Medal in 1992 for work on automatic control, and the Faraday Medal, in 1993, for outstanding contributions to electrical science. He has been Chairman of the Scottish Council for Research in Education since September 1992.

Robin Smith and David G. Griffiths

Distributed Management of Future Global Multi-Service Networks

The current thinking that pervades network and service management is a centralist approach. In this article, the authors develop a philosophy for a fully distributed management system. It is contended that a distributed and cooperative approach to the management of modern telecommunications systems will provide the necessary efficiency, flexibility and capability required. This article is intended as an introduction to this new field of distributed management through the use of software agents.

Introduction

The complexity and operational characteristics of future global multi-service networks (GMSNs) will lead to requirements beyond the capabilities of current centralised management systems. It is contended that a distributed rather than a centralised approach will be capable of meeting these needs. Distributed management systems will require advanced information processing techniques¹ that communicate and cooperate to solve problems. These knowledge-intensive approaches will facilitate the management of GMSNs to provide excellent service while optimising price/performance for the operator. In this article, an insight is given into current research which offers novel ways of managing the complex networks and services of the future.

Background to Multi-Service Networking

A multi-service network (MSN) is defined as any network capable of supporting a wide range of services. The customers for such networks are primarily large corporations with many sites and activities distributed worldwide. To such customers, networks appear to be a private switched facility providing at least the same functionality that they enjoy from international leased circuits. In reality, the service will be supported by a number of diverse networks provided by independent operators. Thus the customer has a virtual network as depicted in Figure 1.

Providing international multi-service capabilities requires considerable capital outlay for infrastructure, interfacing, management and control. To maintain operating costs within tight bounds, extensive automation of

Figure 1 – Virtual network

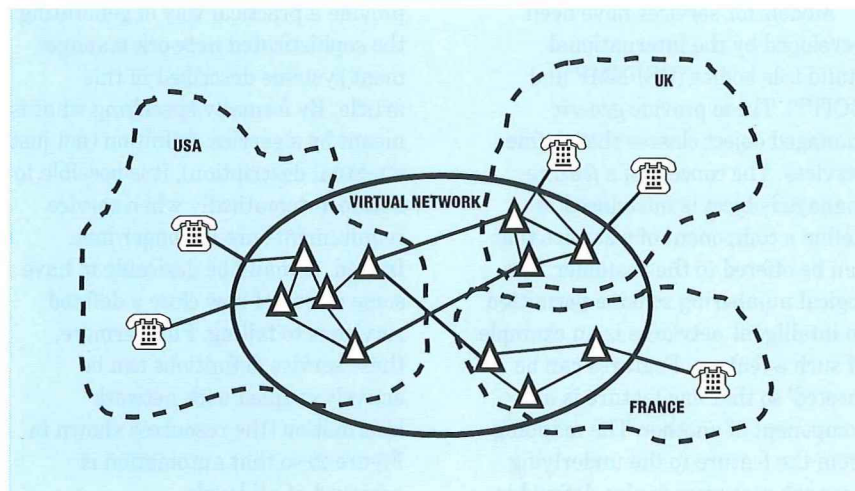


Figure 2—Service-feature-resource dependency

the management and control functions is required.

Service Definitions

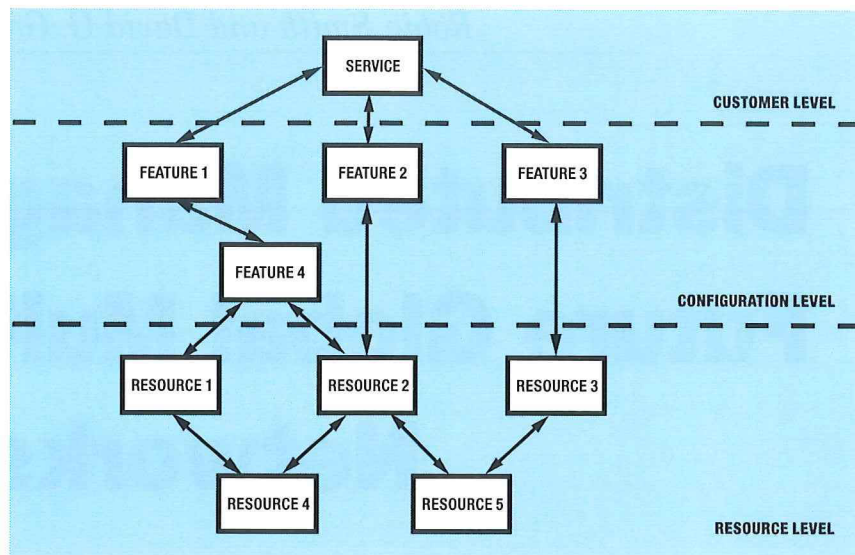
International companies are increasingly entrusting a large proportion of their global telecommunications requirements to one, or at least a small number of, service providers, and it is therefore important they are provided with an excellent level of service. The range of services available is potentially large, evolutionary and dynamic.

Customers also require the ability to manage key features of their own (virtual) network so that services can be requested, altered and ceased from on-line connections to the service provider's equipment.

The definition of a service includes:

- Quality of service (blocking probabilities, bit error rate, error free seconds etc.)
- Target and guaranteed minimum provision times
- Target and guaranteed minimum cessation time
- Target and guaranteed minimum repair times
- Target and guaranteed service availability

Models for services have been developed by the international standards bodies (OSI/NMF and CCITT). These provide *generic* managed object classes that define services. The concept of a *feature* managed object is introduced to define a component of a service that can be offered to the customer. The logical numbering scheme permitted in intelligent networks is an example of such a feature. Features can be 'nested' so that one feature is a component of another. The mapping from the feature to the underlying network resources is also defined in



the feature object. In an intelligent network, the service control point would typically be a resource on which many features (for example, logical numbering, time-of-day routing) depend.

A service is defined in terms of the component features that support the service in question. A service definition will typically refer to a number of features, which in turn may refer to other features and resources. To support this relationship, a number of dependency relationship types are defined (supports, depends_on etc.). Figure 2 shows an example service-feature-resource dependency diagram.

The manner in which customers can be connected to a GMSN is shown in Figure 3. Customer switch access to MSN nodes can be either by direct connection where high bandwidth is required or via intermediate networks such as the public switched telephone network (PSTN).

Distributed Management

Historically, centralised solutions are dominant in all forms of network management and, therefore, only the advantages of distributed solutions are discussed, taking the benefits of a centralised approach as read.

The range of services available is potentially large, evolutionary and dynamic

Formal definitions of service provide a practical way of generating the sophisticated network management systems described in this article. By formally specifying what is meant by a service definition (not just a textual description), it is possible to decide automatically when service requirements are no longer met. Indeed, it would be desirable to have some notion of how close a defined service is to failing. Furthermore, these service definitions can be actively coupled with network information (the resources shown in Figure 2) so that automation is achieved at all levels.

Centralised solutions limit the extent of automation because the practicalities of software engineering quickly dominate as the network control centre grows to accommodate new functionality². Issues of modularity, fault tolerance and scalability indicate that, logically, the functionality should be split into a number of components, each acting largely autonomously³. Figure 4 provides some comparisons of centralised and distributed management from a systems and software viewpoint. There may well be practical benefits of physically splitting the functionality to match

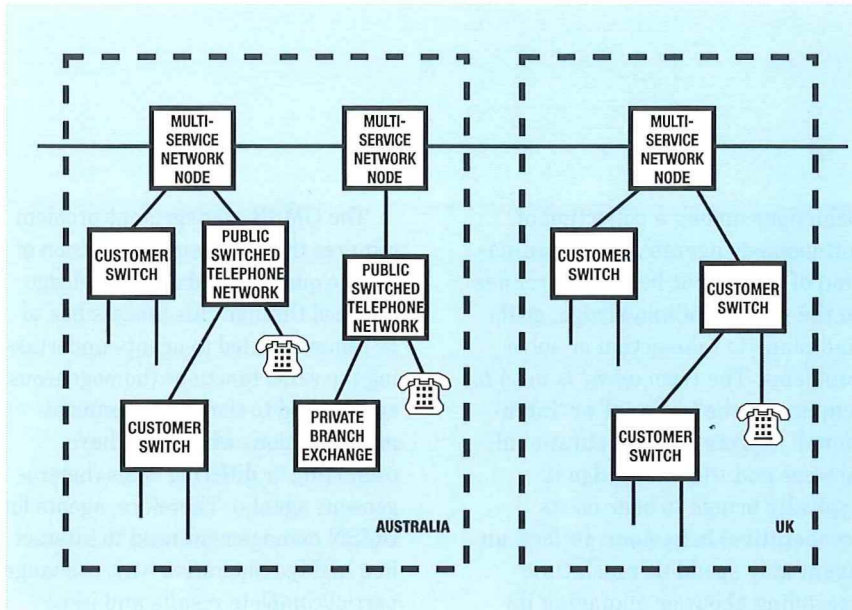


Figure 3—Global multi-service network

analysis, is outweighed by the fact that time taken to respond can be reduced significantly. In future networks, where service and customer demands will be increasingly dynamic, a distributed management approach seems appropriate.

Local knowledge

For networks of moderate size, it is possible for centralised network controllers to have detailed knowledge of the network and use this in detecting and rectifying problems. In a large network, especially one which interacts with many sub-networks, such total knowledge cannot be assumed. Therefore, a system based on local controllers (which have the requisite local knowledge) which cooperate to solve network problems appears intuitively correct.

Less administration traffic

The amount of administration traffic can be large in an environment where routing tables, stored at switches, are subject to change. For example, a customer may specify that when capacity reaches a certain traffic threshold, or at a given time, extra circuits should be made available. To keep track of the network configuration would require the switches to be polled, resulting in a significant and largely unnecessary data transfer. It would be more cost effective to resolve local problems locally.

Limit time zone problems

Good business practice dictates that a customer submitting a request for service should be provided with cost and provision information within the shortest possible time. For an international network, multiple time zones can make this difficult owing to the problem of arranging for the appropriate network managers to be present. Furthermore, service provision typically requires access to a large amount of information, and it is likely that administrative systems are network specific.

In a distributed approach, service provisioning systems can be deployed

the geographical distribution of the network⁴.

The following five facets of system management illustrate how a distributed approach provides overall performance improvements.

Fault tolerance

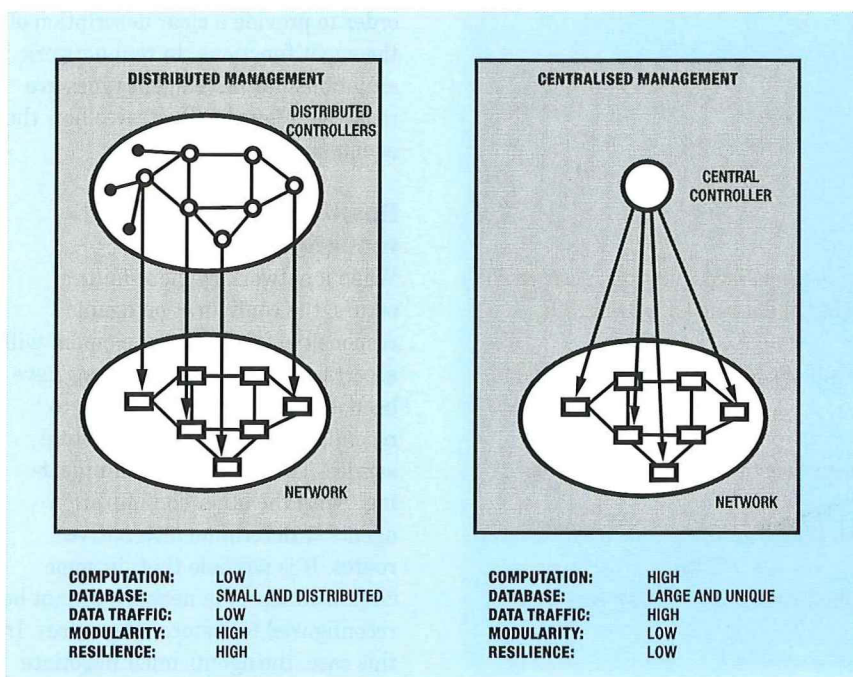
The distributed controllers could be double parented so that one controller could manage a neighbouring switch if necessary. This provides some degree of fault tolerance. Making a centralised network management system fault tolerant is expensive because of the need for multiple

copies of software, error checking of components etc.

Faster response

There are times when a centralised solution appears to perform better. For example, when a large switch or route fails, the best reconfiguration is more easily seen from a global perspective. However, the more common small failures are probably best dealt with locally or with local controllers cooperating. A distributed system will detect problems faster and so the benefit of an optimal reconfiguration, based on global

Figure 4—Comparison of centralised and distributed management



at each node to use local knowledge and thereby reduce the overall provisioning time.

Agent Technology—An Overview

Until recently, artificial intelligence research has, in the main, focused on attributing behaviour to an intelligent process⁵. This behaviour has included problem solving, planning, perception and limited learning. This has proved adequate for a limited class of problems. However, the GMSN management problem is naturally distributed and, therefore, requires knowledge of local conditions. This local knowledge has the potential to be shared by a number of dispersed management systems that can cooperate and negotiate to ensure this knowledge is integrated. Agent theory has evolved to address these issues⁶.

Agent theory does not impose a rigid centralised regime upon cooperation, rather it is concerned with coordinating intelligent

behaviour among a collection of autonomous agents⁷. The coordination of intelligent behaviour relates to the sharing of knowledge, goals and plans to take action or solve problems. The term *agent* is used to emphasise the 'rational' or 'intentional' aspects of a computational process and the knowledge it typically brings to bear on its (cooperative) behaviour. In fact, an agent may spend as much time reasoning about or managing its interaction with other agents as solving its domain-specific problems⁸. This emphasises the importance of multiagent coordination. Having a group of agents does not necessarily mean that they are capable of solving a greater number of problems than an individual agent. Care has to be taken to coordinate the agents so that conflicts and duplication of effort are minimised⁹. A system which achieves this and enables agents to work together to solve a global problem is said to be *globally coherent*¹⁰.

The GMSN management problem requires the processing and fusion of a large quantity of data. Knowledge attained through this process has to be communicated to agents undertaking the same functions (homogeneous agents) and to those with management functions which may have conflicting or differing goals (heterogeneous agents). Therefore, agents for GMSN management need to interact in a highly cooperative way: exchange partial/complete results and issue requests for information in order to satisfy some goal or task. Furthermore, each agent will perform useful processing by using incomplete information while simultaneously exchanging/requesting information/results with other agents in order to construct a complete solution¹¹. Such a system is said to be *functionally accurate* because it exhibits acceptable input/output behaviour. The problem-solving requirement is such that agents need to cooperate to converge on a complete and consistent solution.

Agents in Global Multi-Service Networks

The following GMSN examples illustrate typical agent functions and interactions in a simplified manner in order to provide a clear description of the agent functions. In real network situations, far more agent types are required. Figure 5 illustrates how the agents are interconnected.

Real-time network reconfiguration

When a network element failure occurs, the configuration agent responsible for the failed element will ascertain which, if any, services have been affected. It then attempts to reconfigure the network to restore service. This involves communicating¹² with the other configuration agents to determine alternative routes. It is possible that, in some circumstances, the network cannot be reconfigured to restore all services. In this case, the agents must negotiate

Agent Definition

An agent is a self-contained software program which has two parts:

- a generic part which performs:

communications	}
cooperation	} with other agents
negotiation	}

- a role-specific part which carries out the specialised functions associated with its task. In this article, the following agent roles are described:

customer agent—conducts tasks on behalf of customers using telecommunication services

customer service agent—the network interface to the customer agent

configuration agent—manages network resources to maintain service

service provisioning agent—manages the introduction of new services.

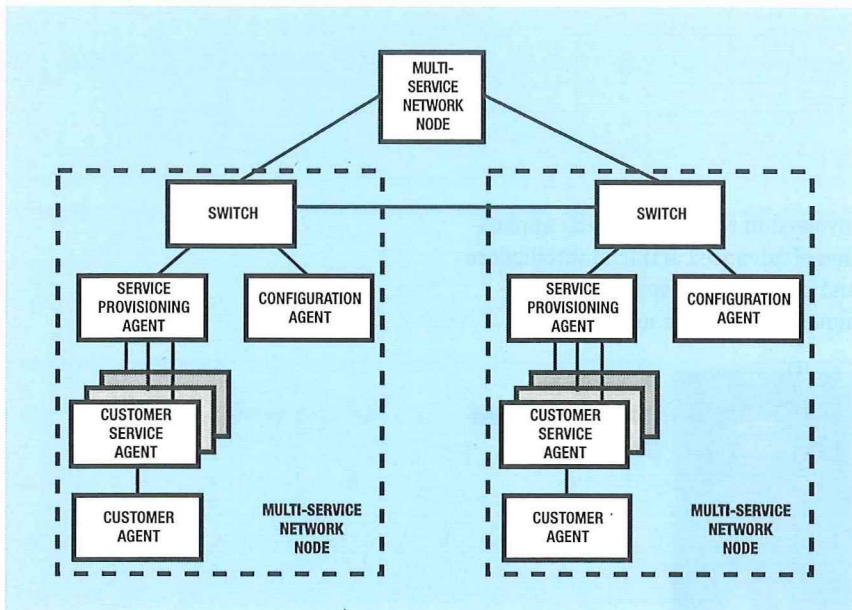


Figure 5—Simple agent configuration

to determine which services should be given priority. The process of prioritising customers' services involves taking note of costs and contracted grade of service etc.

The range of possible configurations is large and an agent may take too long to find a reconfiguration that would restore all services. It may be that a quick but sub-optimal solution would be acceptable. In order to find these quick solutions, the configuration agents will employ shortcuts to guide the process.

These new approaches are considered candidate solutions to the problems created by the rapid growth and complexity of global communications

Service provisioning

When a request is made by a customer agent to provide a connection, the service-provisioning agent commands the appropriate reconfiguration of the network. This may require providing extra equipment and involve requesting the local operator to provide new services at the customer site. In formulating the work schedule, the provisioning agents at each MSN node have to be aware of the provisioning times of the local operator and of the existing configuration.

Long-term considerations such as the need to minimise the number of switches and routes have to be addressed by the agents when performing reconfigurations.

Future Directions

An outline of multiservice networks and a formal description of service definitions has been introduced as a prelude to a discussion of distributed management and agent technology. These new approaches to telecommu-

nications system management are considered candidate solutions to the problems being created by the rapid growth in scale and complexity of global communications. It is believed that agent solutions will facilitate software reuse and thereby help contain the growing software mountain associated with centralised network controllers.

The ideas presented in this article are based on collaboration with the worldwide artificial intelligence research community. Therefore, a sound underlying theory exists, but

there remains a great deal of work to be done before the agent approach will be accepted into mainstream telecommunications management thinking. Areas such as network integrity under agent malfunction and the validation of the decisions made by a community of agents still have to be addressed and studied in depth.

References

- 1 MAMDANI, E. H. and SMITH, R. Advanced Information Processing for Network Management. *Br. Telecom Technol. J.*, July 1991, 9(3).
- 2 Network Administration for the 1990s. *Br. Telecommun. Eng.*, Oct. 1990, 9(3), Special Issue.
- 3 WILLETTS, K. J. Cooperative Management: The Key to Managing Customer Networks. *ibid.*, Oct. 1991, 10(3), p. 174.
- 4 AIDAROUS, S. E.; BALL, T.; TAM, R. and BIGGS, D. A Distributed Architecture for Customer Network Control. *Proc. ICC 88*, Philadelphia, 1988, pp. 11–15.
- 5 WRIGHT, J. R. and VESONDER, G. T. Expert Systems in Telecommunications. *Expert Systems with Applications*, 1990, 1, pp. 127–136.
- 6 GASSER, L. Social Conception of Knowledge and Action: DAI Foundations and Open Systems Semantics. *Artificial Intelligence*, Jan. 1991, 47(1–3).
- 7 DECKER, K. S. Distributed Problem Solving Techniques: A Survey. *IEEE Trans. Syst. Man, and Cybernetics*, Sept./Oct. 1987, SMC-17(5).
- 8 DAVIS, R. and SMITH, R. G. Negotiation as a Metaphor for Distributed Problem Solving. *ibid.*, 1983, 20(1).

- 9 BIGHAM, J.; DAVIES, J. and WHITNEY, C. Towards a Logical Framework for Agent Cooperation. *Br. Telecom Technol. J.*, July 1991, **9**(3).
- 10 GRIFFITHS, D. G. and PUROHIT, B. K. Fundamentals of Distributed Artificial Intelligence, *ibid.*, July 1991, **9**(3).
- 11 HEWITT, C. Open Information Systems Semantics for Distributed Artificial Intelligence. *Artificial Intelligence*, Jan. 1991, **47**(1-3).
- 12 COHEN, P. R. and LEVESQUE, H. J. Rational Interaction as the Basis for Communication, SRI note 433 in COHEN, P., MORGAN, J. and POLLACK, M. (Eds), *Intentions in Communication*, MIT Press, 1988.

Biographies



David Griffiths
BT Development and
Procurement

David Griffiths obtained a B.Sc. in Electronic and Electrical Engineering from King's College London in 1974. He then took up teaching at Brooklands and then Southend Technical Colleges. During this period he obtained an M.Sc. in Dynamic Systems Control. In 1984, he undertook full time PhD. studies in Artificial Intelligence, at Queen Mary College, London. In 1986, he joined BT Laboratories to work on intelligent databases in the knowledge-based systems section. In 1988, he moved to work on a RACE programme that was undertaking investigations into the application of network and customer administration services for integrated broadband communication networks. In 1990, he became leader of the advanced information processing research group. He is currently leading a team

involved in research into the application of advanced artificial intelligence and knowledge-based system techniques for network management.



Robin Smith
BT Development and
Procurement

Robin Smith began his career with BT during 1959, in London, as a technician. After a period working on the installation and maintenance of local area telecommunications equipment, he worked for a number of years on the development of data transmission equipment (modems) when he formed strong links with BT's Research Department. A change in career gave him the responsibility for managing the design of System X, where he headed the central design team. In 1986, he formed the advanced information processing section in the network management department. The section now conducts research into information processing with the Research Department. He is a Chartered Engineer and a Member of the Institution of Electrical Engineers.

Grace Hwang and Paul Seah

Planning the Concert Packet Services Network

Processes adopted in planning today's networks must be responsive to ever-changing market, technological, economic and regulatory forces. This article describes how Concert, BT and MCI plan the Concert Packet Services network, one of the world's largest value-added networks.

Introduction

The Concert Packet Services (CPS) network is the physical network that supports a large number of BT and MCI data services. CPS is part of Concert's portfolio of data services collectively known as *Concert Global Managed Data Services*. The CPS network is a global seamless network jointly operated by Concert, BT, and MCI. The CPS network, which uses TYMNET technology, began operating as a US-based public network in 1971 with 30 nodes. Today, it has grown to become one of the world's largest value-added networks, with over 5000 nodes and 10 000 links in over 100 countries. Every day, the CPS network handles over a million sessions and some 30 000 million characters of terminal-to-host, electronic mail, file transfer, financial transaction and local area network (LAN) traffic.

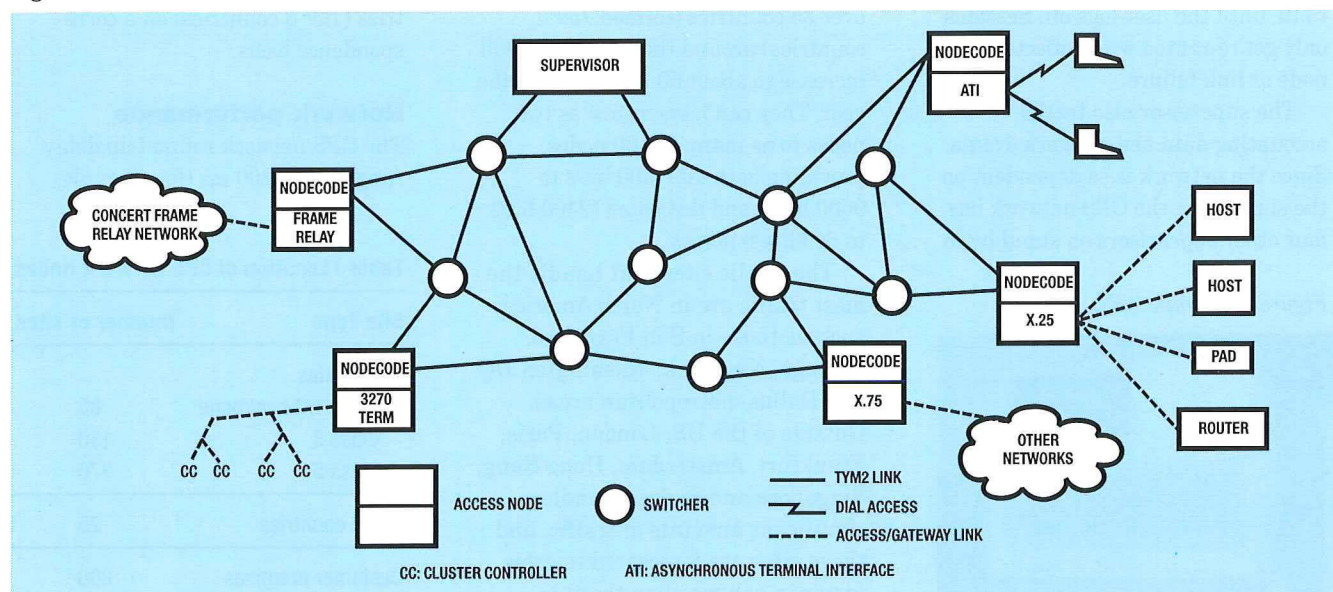
All large networks need to be monitored, analysed, redesigned, and implemented. The sheer size and volume of data on this single network mandates that continual, global and clever planning processes be in place, sensitive to market, technological, economic and regulatory forces. This article describes how Concert, BT and MCI plan the Concert packet services network.

Network Overview and Technology

Access nodes and switchers

There are basically two kinds of nodes in the network, access nodes and switchers, as illustrated in Figure 1. Access nodes interface with end devices such as terminals, modems, hosts, routers, packet assembler/dissassemblers (PADs) communicating using protocols such as X.25, 3270

Figure 1—CPS network overview



SNA, SDLC, UTS, etc. These external protocols are converted into an internal X.25-like protocol, called the *Tym2* protocol, and forwarded to switchers as packets. The nodal software that handles Tym2 packets is called *NodeCode*.

A single Tym2 packet varies from 16 bytes to 128 bytes, and could contain data from multiple users. On average, about 25% of a packet is protocol overhead.

The role of the switchers is to deliver Tym2 packets to the intended destination device, usually via other intermediate nodes, in a store-and-forward fashion. At each hop, error-checking is done; Tym2 links use a sliding window mechanism to recover flawed packets.

Supervisor

The mechanics of routing and session establishment are performed by a specialised node in the network called the *supervisor* node. This node has visibility to all nodes and links in the CPS network. Each time a user logs into the network, the supervisor establishes a virtual circuit that is primarily based on least hops. Routing decisions also take into account link speeds, network load, whether a session is batch or interactive, and whether a link is terrestrial or satellite. Once the route is determined, the session remains on that route until the user logs off. Sessions only get re-routed when affected by a node or link failure.

The supervisor also tracks accounting data and network events. Since the network is so dependent on the supervisor, the CPS network has four other supervisors on stand-by in

case the active one fails. The five supervisors take turns being the active supervisor; the takeover process occurs every 24 hours.

Hardware

Basically, two families of nodes are currently deployed in the CPS network. The Engine family, which is based on proprietary architecture, currently has two models: the PXL for desktop switching, and the CXL for mid-range to high-end switching (see Figure 2). Older models of the Engine nodes such as the Full Engine, Mini Engine, Pico Engines still exist in the network, but they are slowly being phased out.

The Turbo Engine family, whose hardware is more modular and based on an open architecture, is designed for mid-range to high-end switching.

CPS Network Today

Public sites

CPS network nodes are located in over 1400 sites throughout the world. Most of the nodes are in so-called *public* sites, where BT or MCI owns, installs, and manages the equipment. In some countries, such as Austria or South Africa, nodes in public sites are operated through third-party agreements.

Today, public sites can be found in over 35 countries (termed *Tier 1* countries) around the world, and will increase to about 50 by the end of the year. They can have as few as two nodes to as many as 100 nodes, providing both dial (300 bit/s to 9600 bit/s) and dedicated (2400 bit/s to 64 kbit/s) access.

The public sites that handle the most traffic are in North America, such as those in San Francisco, Chicago, New York, Washington DC and Dallas metropolitan areas. Outside of the US, London, Paris, Frankfurt, Amsterdam, Hong Kong, Singapore and Sydney handle significant amounts of traffic, and these sites are forecasted to grow faster in volume than those in

North America. These major sites constitute *Class 1* (primary backbone) and *Class 2* (secondary backbone) sites.

Public sites are also in other significant population centres—for example, Marseille (France), Calgary (Canada)—that provide tertiary concentration; these sites are categorised as *Class 3* sites. And finally, there are *Class 5* sites, smallest of the public sites that provide only asynchronous dial access to local communities. See Table 1.

CPS network nodes may also be found in many large customer premises, some forming substantial subnets.

Other sites, gateways

In over 25 countries (*Tier 2* countries), CPS network nodes are owned and operated by other network service providers. Some providers such as Network Information Service (NIS) in Japan even have their own separate TYMNET network that is gatewayed to the CPS network. Since the CPS network supervisor has no view to other TYMNET networks, users from those countries do not receive the full end-to-end support, centralised billing or universal login that users in Tier 1 countries have.

The CPS network also has X.75 gateways set up with over 50 countries (*Tier 3* countries) on a correspondence basis.

Network performance

The CPS network round trip delay ranges from 200 ms (for example,

Table 1 Location of CPS network nodes

Site Type	Number of sites
Public sites	
Class 1/2 backbone	80
Class 3	150
Class 5	370
Tier 2 countries	25
Customer premises	800

Figure 2—Compact XL



Figure 3—CPS network functional hierarchy

within the UK) to about 1.2 s (for example, Asia to Europe). The hierarchical structure keeps the number of network hops generally between three and seven hops.

Network availability varies from 99.99% in the well-meshed backbone to about 99.7% for users with dedicated ports, and to about 99% for dial access users.

CPS Network Architecture

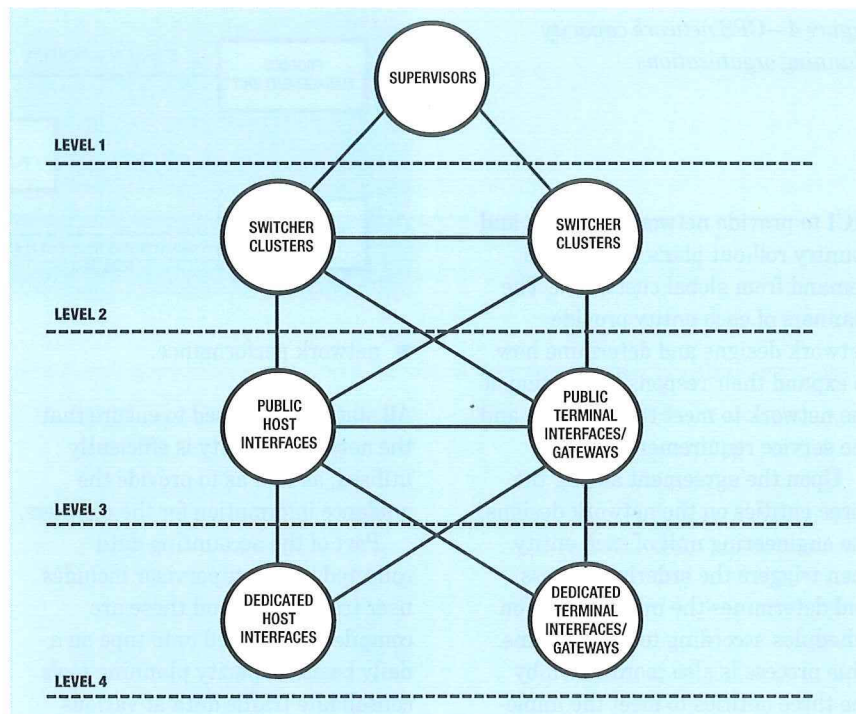
TYMNET-protocol links in the public network are mostly comprised of speeds ranging from 9.6 kbit/s to 64 kbit/s. Class 5 sites are generally connected to Class 3 or backbone sites. Class 3 sites are connected to other Class 3 sites and backbone sites. Backbone sites are richly connected to other backbone sites in the same continent; however, connectivity to backbone sites in other continents tends to be limited to a few key sites.

This relationship describes a geographical hierarchy. Traffic from numerous customers and small public sites is consolidated towards the fewer backbone sites, where most of the switching takes place.

Most of the backbone links are carried on a worldwide time-division multiplexed (TDM) platform, which is managed by Concert's global managed platform (GMP). This platform is based on Network Equipment Technology's (NET) IDNX technology using E1 and T1 trunks. A future article will describe the GMP in greater detail.

While most public network links are digital and terrestrial, the CPS network has a small number of international links that are satellite. These are usually links to Tier 2 countries or backup links between Tier 1 countries.

The CPS network is also organised by another kind of hierarchy, one that is based on function. See Figure 3. Access nodes comprise the two lower levels of the hierarchy. Nodes in level 4 are customer premise equipment; nodes in level 3 are public site



nodes that directly interface with terminals, hosts and other networks.

Level 2 nodes switch only pass-through traffic. Backbone sites have highly meshed switchers that switch packets from other sites as well as packets from local access nodes. Where a site contains a supervisor, that supervisor is connected to the local switcher cluster. Through this functional hierarchy, the supervisor can monitor and control all nodes in the CPS network.

Capacity Planning Objectives

In order to expand the CPS network to meet the growing demand in a cost-effective manner, the network capacity needs to be carefully planned based on the statistical analysis of the network and market forecasts.

Three objectives are associated with capacity planning:

- to ensure sufficient network capacity to meet customers demand within a 6–12 month planning horizon,
- to meet the performance criteria and deliver quality of service to customers, and
- to best utilise the existing network resources and adjust the network capacity in a cost-effective manner.

To meet these objectives, some trade-off needs to be made in the network design since all three objectives can be conflicting. For example, to maximise the performance level, it will be ideal to connect customers directly to their destinations, which will result in a one-hop delay versus a multiple-hop delay. However, such a solution is too costly to operate, and it would not be feasible to offer competitive pricing to customers.

Owing to the complexity and the size of the CPS network, heuristic network designs which meet all three objectives are delivered periodically to expand the network capacity.

Capacity Planning Organisation

Capacity planning of the CPS network is a joint effort of the three entities: Concert, BT and MCI. Concert delivers global network designs for transborder connectivity. BT and MCI deliver in-country network designs for their responsible territories. The pool of expertise from all three entities, each specialising in planning at the country, regional or transregional level, is what it takes to ensure that the global network operates seamlessly and cost-effectively.

The product management unit provides the service requirements that are agreed among the three entities. The product marketing unit in Concert coordinates with BT and

Figure 4—CPS network capacity planning organisations

MCI to provide network forecasts and country roll-out plans to meet the demand from global customers. The planners of each entity provide network designs and determine how to expand their responsible portion of the network to meet the forecasts and the service requirements.

Upon the agreement among the three entities on the network designs, the engineering unit of each entity then triggers the ordering process and determines the implementation schedules according to design plans. This process is also coordinated by the three entities to meet the implementation schedules. Figure 4 outlines the relationships of the planning organisations.

Elements Of Capacity Planning

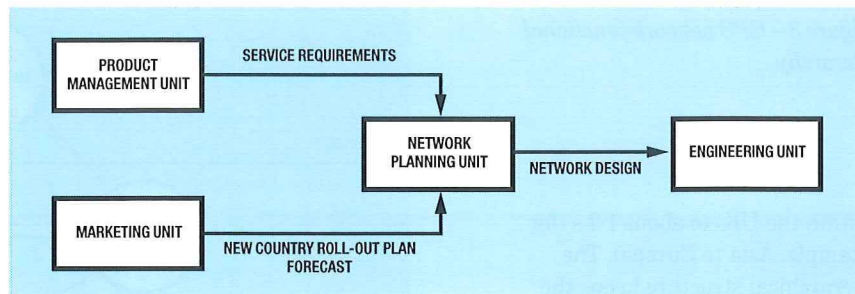
The four main elements of CPS network capacity planning are shown in Figure 5. They are listed below and described in the following sections.

- gather network statistics on traffic patterns, network performance, and network resources;
- perform network analysis based on the network statistics gathered from monitoring;
- deliver network designs which incorporate network analysis results, forecasts, new country roll-out plans, and service requirements; and
- implement the network design by allocating available capacity from the existing network and expanding the network capacity to meet the demand.

Network statistics collection

Three kinds of network statistics are gathered periodically to monitor closely the network capacity:

- traffic,
- network resources utilisation, and



● network performance.

All statistics are used to ensure that the network capacity is efficiently utilised, as well as to provide the guidance information for the planners.

Part of the accounting data collected by the supervisor includes user traffic data, and these are compiled and stored onto tape on a daily basis. Capacity planning tools consolidate traffic data at various hierarchical levels (host, node, site, class, country level, and network-wide) and time intervals (by hour-of-day, day-of-week, month). These higher-level statistics paint a general usage pattern, the hourly and daily peaks, as well as the rate of traffic change in the CPS network.

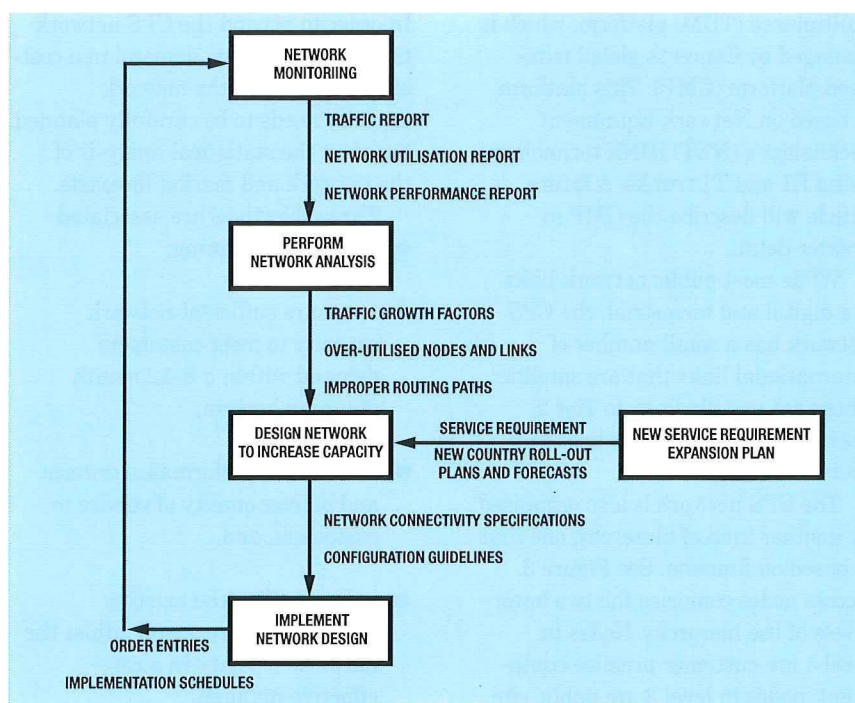
Network resources are closely monitored by measuring the network utilisation and keeping track of the evolving trend of the network. Such network trend statistics are used to identify the capacity within the existing network topology.

Network performance is monitored by several internal tools. DeMon, the delay tool, measures the round-trip delay between nodes. This tool takes a set of nodes as the input and executes the commands within the NodeCode software to send out test packets from the input nodes onto all the active circuits. XRAY and Deluti, the link utilisation tools, measure the utilisation of infrastructure links, not only on an average basis, but also instantaneous peaks. Such measurement is taken by using NodeCode software to take snap-shots of the network and to collect real-time statistics. TRMS, the dial-up access utilisation tool, calculates the congestion level of access sites. The congestion level is computed by using queuing theory to model the number of calls blocked. All tools utilise the data-collection capability within the NodeCode software.

Performing network analysis

The traffic patterns and trends can be observed by analysing the traffic

Figure 5—Main elements of CPS network capacity planning



statistics. The traffic patterns reveal traffic volumes at top sites as well as flows to their favoured destinations. Since the planning horizon is more than six months, forecasting has to be done based on historical trend. The traffic trend, which shows the traffic fluctuation, is translated into overall as well as regional traffic growth factors. Both traffic patterns and growth factors are used to determine architecture upgrade, site reclassification, and backbone infrastructure redesign.

Network analysis is performed based on the results of the above-mentioned network statistics. By analysing the utilisation statistics at the node, link and access port levels, the planners can observe whether network resources are adequate to support the demand. At a certain threshold, the network upgrade process is then triggered.

The delay statistic provides feedbacks on the overall performance of the network. By analysing the routing paths that have caused unacceptable delays above the targeted threshold, the planners determine whether the problem is caused by a shortage of network capacity or poor connectivity of switchers at the site level. The process of adding capacity or redesigning site clusters is then triggered according to the result of the analysis.

Designing the network

Network capacity is assigned and added according to the network connectivity specifications provided in network design plans. The designing process is triggered by a few factors:

- large customer bids;
- performance results do not meet the internal threshold;
- new service requirements provided by product management unit;
- new country roll-out plans provided by marketing; and

- periodic forecasts based on the analysis of traffic trend and marketing data.

The first three are ad hoc factors that trigger the need for network re-engineering, while the last two factors drive periodic (three or four time a year) network audit and design.

Based on the analysis of network utilisation, the planners add network capacity and improve the infrastructure on an ad hoc basis when network utilisation crosses certain thresholds at the node and link levels. New nodes and links are added to the existing network topology according to the functional hierarchy in Figure 3. When a new service is proposed to be added to the Concert packet services portfolio, the planners provide network designs to increase network capacity according to the service requirements and the service roll-out plans.

Backbone design

Based on the forecast and the new country roll-out plans provided by the marketing unit periodically, the planners expand the network infrastructure to meet such requirements. When expanding the network infrastructure based on the forecast, the planners grow the traffic based on the predicted growth factors on a node level and then consolidate the traffic into the site level. In order to provide sufficient capacity to accommodate traffic during the peak time, the planners size the bandwidth requirement based on the peak traffic instead of the average traffic. A set of bandwidth sizing tools, called *Configurator*, which model the supervisor least-hop routing algorithm, are used to automate the design process.

Access design

The expansion of dial-up access capacity is performed in a similar fashion. A set of tools, which grow the traffic based on forecast and size the dial-up port count to be added, are used to engineer capacity and

hardware resources to meet the customers' demand and maintain the quality of service.

Site design

The new nodes and links identified in the backbone and access design process are added to a site according to a set of rules specified in the guidelines. At a large site, such as a Class 1 or Class 2 site, the nodes are configured and connected according to the CPS network functional hierarchy as specified in Figure 3. At a small site, such as a Class 3 or Class 5 site, the nodes are configured and connected in a hybrid environment in which a node is configured for both access and switching functions.

Implementing the network design

This process implements the network design by ordering appropriate hardware equipment and/or telco circuits to meet demand forecasts. The ordering process is performed by using the internal database and ordering entry system known as *order management and network information* (OMNI). This system keeps track of all the modifications on the network topology as well as the existing network capacity.

With the information provided by OMNI, planners identify the existing hardware capacity and assign the available capacity to meet the network design. When there is insufficient capacity, additional supporting equipment is ordered through OMNI.

All the order entries related to the network capacity expansion due to traffic growth are performed based on just-in-time philosophy. However, other order entries related to new country roll-out and new service requirements are performed based on the targeted schedules defined by the marketing unit.

Upon the completion of the order entry process, the implementation unit then defines the project implementation schedules and coordinates with field engineers to install the necessary equipment and circuits.

Future Development of the CPS Network

With increasing competition in the global networking market, it is necessary constantly to improve quality while reducing costs. Four plans are targeted to improve the quality of the CPS:

- increase backbone bandwidth,
- provide high-speed 14.4 kbit/s dial-up access,
- deploy new hardware platform, and
- establish a point-to-point protocol (PPP) gateway service.

With the rapid growth of customer demand, it is more cost-effective to consolidate low-speed links into high-speed links. Such bandwidth consolidation not only can reduce hardware cost but also can improve network performance. The high-speed service roll-out plan is under study by the design groups of Concert, BT and MCI.

With the development of new network technology and the growing demand for network access, more dial-up users are requesting the high speed dial-up service for faster transmission rates. Currently, the 14.4 kbit/s dial-up service is set up within the CPS network for beta test. The specific roll-out plan is yet to be defined by Concert, BT and MCI.

With the growing demand for the high-speed infrastructure and the high-speed dial-up service, a new hardware platform is proposed to support the high-speed connectivity. By implementing a new hardware platform with higher switching capability and port density, the network operation cost can be significantly reduced while improving the quality of service. In addition, with the new hardware platform, the CPS network can utilise switched multi-megabit data service (SMDS) and asynchronous transfer mode (ATM) technologies as transport

mechanisms in the future. Currently, the development unit is working closely with prospective vendors to provide system evaluations.

The PPP gateway service provides the high-speed access by using the CPS network dial-up access service and the frame relay backbone infrastructure. This service allows dial-up users to transmit data at a higher speed by converting Tym2 packets into frame relay cells. Currently, the service descriptions and requirements have been defined. The development unit is testing system integration in the laboratory.

Conclusion

With the growing customer demand for global connectivity, the CPS network continues to expand to new cities and countries. New countries pose new challenges as each country has unique regulation, availability of data facility, lead time, market demand, and tariffs. Thus, capacity planning processes will continue to be enhanced in order to accommodate the diversity of new countries.

Biographies



Grace Hwang
Concert

Grace Hwang joined BT North America Inc in 1992 after obtaining an MS degree in Operations Research from Stanford University. As Senior Network Design Engineer in Network Planning Department, she was responsible for providing network design and capacity planning for the Tymnet global network. With the emphasis in supporting Tymnet packet services, she was involved in customer bid design, network performance analysis for service level agreements, defining

guidelines and rules for network planning, and expanding Tymnet infrastructures into Asia Pacific. In her current role in Concert, she continues to have the planning and design authority for the CPS network. In addition, she is responsible for providing business cases to request financial approval for expanding the CPS network. She is currently working on expanding the CPS network from 37 countries to 63 countries, rolling out a new hardware platform for the CPS network and interconnecting two Concert-owned frame relay platforms.



Paul Seah
Concert

Paul Seah joined McDonnell Douglas Network Systems Company in San Jose in 1989, after obtaining an MSEE degree from Stanford University. This business unit subsequently became BT North America in 1990. As Senior Network Analyst, he was responsible for planning and designing the Tymnet public network. He was also actively involved in major customer bids, handling issues on network performance, provisioning and operations. He was the principal designer of BT's global frame relay network when the GNS LAN Interconnect service was launched in 1992 in 15 countries. In his present role in Concert, he continues to be the planning and design authority for the Concert frame relay network. His current projects include design plans to grow the network to over 30 countries, to upgrade the network to support network integration with MCI's Hyperstream network, and to re-architect the network to multiple domains.

Paul Warren

Ensuring Quality in Design of Customer Networks

BT's continued global presence in the customer networks arena is dependent on the increasing use of automated processes enabling more responsive, robust designs which are most appropriate to individual customer needs. These must be augmented by integrated resource management, project management, and skills databases which will allow customer bids to be addressed quickly by people with the right skills.

Introduction

Since BT was privatised some ten years ago, regulation, competition, and the need to expand into overseas markets have resulted in BT becoming a key player in the global arena. BT has offices in Australia, Japan, Singapore, most Western European countries, as well as a presence in North America, where BT's interests are now represented by Concert, the BT/MCI subsidiary.

Historically, BT's international business has been *correspondent*, where BT has been responsible for the UK 'half-end' of an international circuit, and the overseas PTT is responsible for the other half. Over the last couple of years, there has been a move towards *non-correspondent* services, where BT actually provides the whole end-to-end service. This has obvious advantages in that half the income from the service does not have to be given over to the other PTT, but competition is fierce with other carriers also vying for the same business.

In the case of both correspondent and non-correspondent working, carriers effectively have access to the same technology; this means that any superiority in technology is only a short-term advantage; the long-term differentiator for the potential customer is in the quality of service provided and the assurance that the customer is getting the network solution best suited to needs. To gain the advantage and win business, BT must provide this edge at little or no extra, and preferably at less, cost.

Global Networks

Nowadays, many companies are global, having offices, factories, and

warehouses in different parts of the world. These customers want integrated seamless communications between these areas; they do not want to negotiate with different carriers for each private circuit or network node. BT is in the business of providing a complete solution which will give the customer a single point of contact for service management of the whole network.

Global Networks is a division of Worldwide Networks responsible for the planning, provision, and operation of overseas network infrastructure. Much of the work is undertaken in response to requests from BT Global Communications which is the interface with external customers and determines their requirements.

Global Networks operates in three major zones, Asia/Pacific, Europe, and UK. (Since the UK itself is often a major hub of a company's international network, Global Networks also has a technical, as well as administrative, presence in the UK.) These divisions are administrative rather than functional, since efficiency dictates that resource is used where necessary to even out demand, and obviously a network may span two or even all three of the zones.

Network Planning Processes

A customer that approaches BT (or any other carrier) often does not know what services are available or the technicalities involved; however, the customer has a problem that needs a solution at the service level. A simple process that might follow a customer's enquiry is shown in Figure 1.

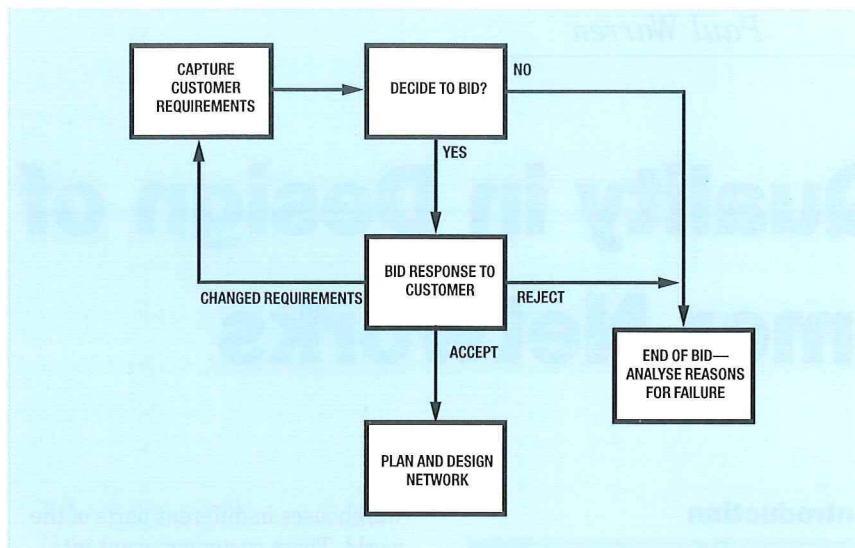


Figure 1—Simplified process for customer network design

BT captures the customer's requirements from this service viewpoint and, after some initial work to determine the feasibility and profitability, puts a bid response to the customer. The customer may accept or reject the bid, or may (as the customer is always right) decide to change requirements, perhaps necessitating a new response.

In all but the first case, the work to date may have been wasted (although lessons can be learned even from unsuccessful bids). When it is considered that putting together a bid may cost between £3M and £10M, and that (depending on the type of bid and service) competition is intense so the win/bid ratio is not high, it can easily be seen that reduction of bid costs and improvements in the win ratio are of primary importance.

Network Design Standards and Tools Unit

Network Design Standards and Tools is a unit within the Regional Network Planning Department of Global Networks, and has responsibility for identifying and delivering relevant network design processes, skills and tools to maintain BT's position as a key global player.

Global Networks is not the sole centre of network design within BT, others being located within Global Communications and National Business Communications Divisions. However, although primarily serving Global Networks' design capability, Network Design Standards and Tools is intended to be the focal point for all

BT's customer network design tools, with the objective that duplication is avoided, and the benefits of a tool or process are available to BT as a whole.

The three areas of concern for Network Design Standards and Tools are skills, processes, and tools. This article will cover tools in some depth, and the others more briefly.

Skills

It has been recognised that competent network designers are a valuable asset which should be nurtured and protected. Although it might appear that the increasing use of automated tools will lessen the need for qualified people, this is not the case. The primary reason for developing and using tools is to speed the design process, freeing the skilled designer from the mundane 'number-crunching' work so he or she can concentrate on the tasks that a tool cannot do, such as creating more imaginative options for the tool to work on.

Skills Database

Every time a bid document needs to be created, or a network is to be worked on, a project manager is assigned. The project manager has to create a team of people with the skills necessary to do the job. Not only must the project manager collect technical skills of the right discipline (X.25, network management, switched multi-megabit data service (SMDS), Cisco router configuration, etc), but other talents may also be an advantage (a foreign language, prior travel to a foreign country, internationally-recognised qualifications, etc).

One of the main people-related problems in a multi-divisional organisation such as BT, is the difficulty in identifying exactly where the designer resource is located, its skill level, product/supplier knowledge, and its availability. Regional Network Planning is in the process of setting up a pilot trial of a skills database, which will hold information on its network designers. A project manager will be able to interrogate the database to ensure that the right people are gathered into the team. Other detail such as the designers' line managers are also held, so the project manager can negotiate priorities where availability is not immediate.

The database will be updated as designers learn new skills, or gain more experience in existing ones. Curricula vitae (CVs) are also held on the database.

Several other divisions within BT are interested in the database concept, and it is possible that the scope could be widened to include all of BT's network design people.

Within Regional Network Planning, the skills database will also interface to other databases for resource and project tracking, and budgetary databases.

Skills Accreditation

A second area of concern is that there is no recognised qualification or accreditation for network designers. This is not a problem internally within a network design company if the designer is known to be competent; however, in the era of quality management, where the demonstration of competence is becoming increasingly sought by potential customers, a skills accreditation scheme is being planned. Worldwide Networks Personnel has taken responsibility for defining and implementing the scheme, although the original concept was that of Network Design Standards and Tools, which is still driving the overall programme.

The intention is to ensure that designers are accredited to defined

Figure 2—Tool areas

standards; these standards will be acceptable to customers as proof that their networks are being designed or maintained by suitably qualified people.

Workshops and Seminars

In a physically dispersed company such as BT, a major problem could be the lack of communication between groups of designers. For example, if a customer approaches BT Managed Network Services (MNS), the solution is likely to be X.25, simply because MNS is a centre of expertise for X.25 within BT. Similarly, approaches to other units may yield solutions dependent on the main expertise of that unit. One of the main aims of Network Design Standards and Tools is to ensure a consistent response to a customer no matter 'which door the customer knocks on'. This means that the customer should be offered a solution that best fits the customer's needs; this implies that network designers, although expert in their own area, should have at least a superficial knowledge of the capabilities and limitations, the advantages and disadvantages, of other technologies and services. To address this, several workshops have already been held at which designers from varying disciplines discuss and learn about related areas.

Processes

Network Design Standards and Tools has visited several design units within BT to document the local processes associated with network design; the intention is to determine whether a single process could be put in place across BT which would ensure the consistent customer response which has already been mentioned. In practice, a single process is improbable considering the disparate work environments, technologies, etc to be found within BT globally; however, some rationalisation and commonality are likely.

The objective of ensuring that network designers have visibility of other services and technologies has

led to the development of the *Network Design Guide*. This is a handbook that contains a section on each available service such as X.25, flexible bandwidth service (FBS), SMDS, etc. The *Network Design Guide* is intended to be a reference for designers which will enable them to be aware of other technologies and to give them a base knowledge of design criteria within these areas.


Within Regional Network Planning, these processes are an integral part of the ISO 9001 approval processes which govern project registration, bid support, and network design activities.

Tools

The scope of network design tools is often defined to be those tools that are used in the planning and design phases; however, it can be argued that any tools which have an input to these phases should be included. Also, since the Regional Network Planning Department of Global Networks works with other Global Networks' units to plan the provision and operation of networks, it can be seen that the scope is not clearly defined; tools such as those used for bid support, pricing, bandwidth management (perhaps leading to inventory management), cable wiring tools, graphical information systems and tariff databases may possibly be included.

If the scope of network design tools can be considered as one dimension of Network Design Standards and Tools' remit, another two dimensions are the services covered (X.25, SMDS, voice, etc), and the 'level':

- service (parameters to be considered are those which are customer-affecting; for example, grade of service, up/downtime);
- network (the parameters are network-affecting; for example, channel quantities, bandwidths); or
- element (for example, detailed configuration of equipment, routing).



This is illustrated in Figure 2.

The intention is that the resulting three-dimensional table should have as many 'cells' populated as is economically and practically possible; it would not be possible to populate every cell since some are mutually exclusive and others overlap.

Tools Portfolio

The Network Design Standards and Tools unit has the responsibility of identifying tasks within the day-to-day work of Global Networks which could be more efficiently carried out by automating part or all of the task. Although the unit tries to be proactive in this role, reliance must be placed on the potential customer being able to recognise those areas where tools would be of most use. In turn, however, for this to happen, the customer must be aware of what tools are already available, what is possible in terms of tool functionality, and, above all, the existence of the Network Design Standards and Tools unit. Tool identification is a mandatory part of the Regional Network Planning key processes. It is seen as an essential element of the unit's programme to maintain consistently high levels of quality and efficiency.

Once a customer has identified an overall requirement, the unit assists in capturing as much detail as possible so that it can be decided (if it is a suitable subject for the unit's interests) whether a tool needs to be developed or whether one is available already. Also, at this stage, an attempt is made to identify other possible customers who may have the same or similar needs.

If a suitable tool is available 'off the shelf', the customer will be pointed in the right direction. It is not the unit's responsibility to purchase proprietary tools, although it will

Figure 3—Example of set of design-related tools

assist in negotiating with suppliers of proprietary tools to obtain quantity discounts, and advantageous contracts. If a tool needs to be developed, the unit liaises with BT Development and Procurement (D&P) for the production of suitable software. This may be by development in-house at one of BT's several software centres, or by negotiating enhancements to proprietary tools.

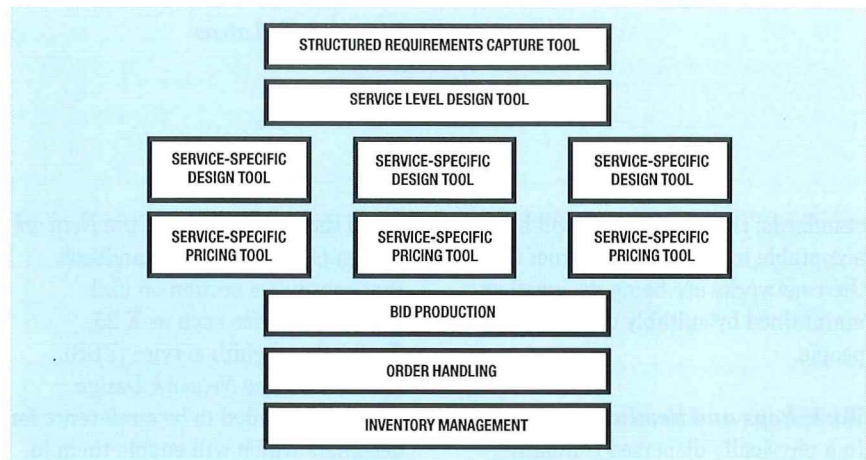
Development of the tool is to agreed time-scales and costs, and is delivered to the customer for acceptance. First-line support for tools is provided by a support desk which is the customers' first point of contact for problems, queries, training enquiries, requests for documentation or software, etc. Second-line support is arranged with the D&P development teams or outside suppliers as appropriate, depending on the tool's source.

To ensure standardisation and compatibility, tools are specified, wherever possible, to run either on IBM PC compatibles or a Sun Unix workstation. Since these are universally available, there are very few proprietary tools that cannot be run on one of these.

The tools portfolio is a list of tools recommended by the Network Design Standards and Tools unit. It is not a comprehensive list, since not all the 'cells' of the three-dimensional table previously mentioned have been populated; even had they been, the pace of new technology and the availability of new services means that the table is dynamic anyway.

The unit also issues a list of known tools. This includes not only those in the portfolio, but also tools which have been evaluated and found wanting or unsuitable, those which have not yet been evaluated, and others which are suitable for purpose but may be of minority interest. Each tool is given a rating for various features such as interfaces, useability, accuracy, upgradeability, etc.

Figure 3 shows a possible set of tools which might be available to designers to enable them to achieve



an end-to-end service. This tool set is not comprehensive, nor are all the tools necessarily associated specifically with network design. Inventory management for example is one area that can encompass more than network inventory, and is not a necessity for design work, but a suitable tool which can track available bandwidth, card slots and hardware in each location can immensely increase the power of automated network design overall by ensuring best use of existing equipment rather than providing new hardware.

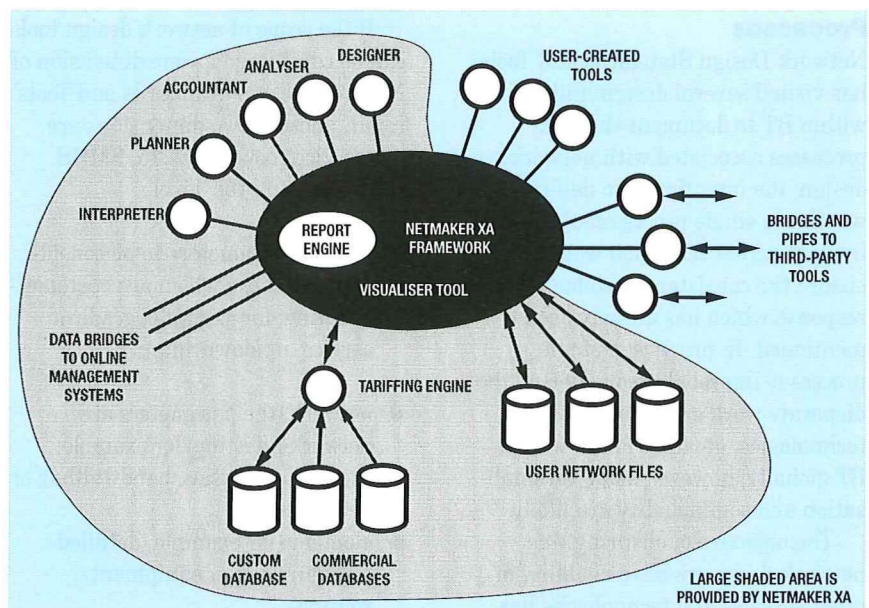
Similarly, interfaces between tools are deliberately not shown in this figure, but the inclusion of these will again increase efficiency because of reduced rekeying of data, transcription errors, etc.

The following sections detail a few of the more common tools in use within BT.

Make NetMaker

Make is a US company which is internationally known for its leading-edge network design tools. NetMaker is one of the company's products which allows design and costing of networks built from various suppliers' equipment. NetMaker XA is the latest version of the tool offering functionality in several network design areas. The tool is modular consisting of core software, around which the user can put other modules depending on the application. Device libraries containing details of manufacturers' router or multiplexer equipment are available to suit the users' requirements; libraries for Timeplex, Cisco, NET and Wellfleet are among those offered. A Sun Sparc 10 is necessary to run XA. Figure 4 shows the interconnection of which NetMaker XA is capable, not only between its own modules, but also with other tools which can be

Figure 4—NetMaker XA



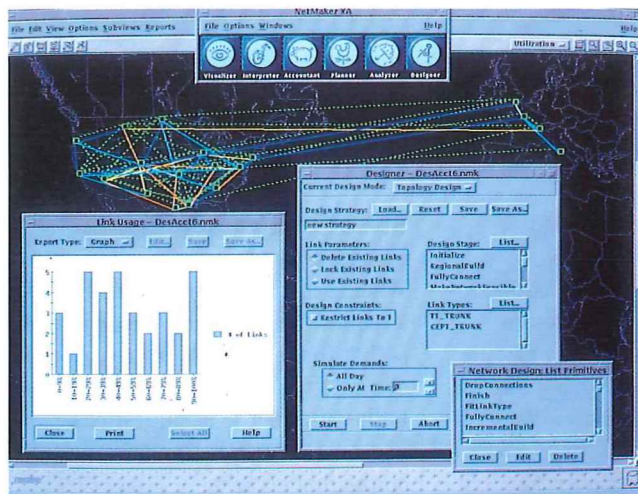


Figure 5—Example of graphical reporting

interfaced so that XA's powerful features can effectively be used to enhance those of the connected tool.

Figures 5 and 6 show typical screen shots of XA. Both figures illustrate the integration of the modules by the row of icons at the top; Figure 5 gives an example of graphical reporting (in this case of link usage for the network), and Figure 6 shows a cluster analysis. Although running under a Unix operating system, the tool interfaces have been designed so that the network designer does not also have to be a Unix expert. The tool displays links as different colours depending on the (percentage bandwidth) utilisation; this is a method used by other tools, but NetMaker also allows the operator to specify the percentage range for each colour. This makes for a very efficient system whereby minor variations in utilisation can be spotted instantly.

Traffic demand can be input to the tool either by manual keying, or by automatic download from 'sniffers'—probes which can be inserted into the network to analyse and store traffic. XA allows the overall traffic to be broken into different applications (for example, Microsoft Mail). The tool can then estimate average extra traffic which would be generated in scenarios such as extra nodes. Optimisation, simulation, survivability analysis and cost tracking are just a few of this tool's features.

NEMO

NEMO is a BT in-house development from the BT Laboratories at Martlesham Heath, Suffolk, England. It is really a platform for several

different tools, some of which can form an integrated toolset, others being stand-alone modules for specific applications. NEMO runs on a Sun Sparc workstation and is extensively used by network designers and systems engineers across BT.

The basis of NEMO is a powerful set of three tool modules that enable private circuit networks to be designed and costed. The modules are:

- **BDT (backbone design tool)**—for designing backbone networks. The tool can automatically suggest topologies and dimension links according to traffic profiles. The resulting backbone can be optimised to give a defined level of resilience at minimum cost.
- **ADT (access design tool)**—for designing access networks. It can model both time-division multiplex (TDM) and Statmux equipment, and designs the access network to achieve connectivity to the backbone at minimal cost. Again, it automatically produces topologies and dimensions links.
- **PACT (pricing and costing tool)**—this is fully integrated with ADT and BDT. It allows storage of customer details, contains full exchange and exchange service details and provides pricing for networks designed on ADT and BDT.

In addition, there are other modules:

- **FDT (Featurenet design tool)**—this is a fully integrated module within

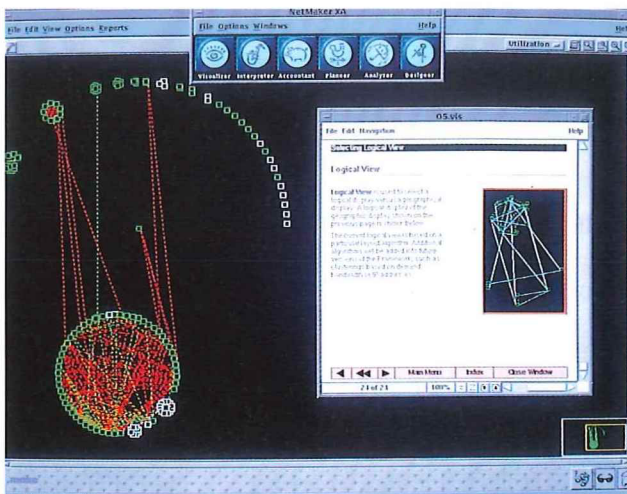


Figure 6—Cluster analysis

the NEMO toolset, allowing the designer to design and cost a complete Featurenet network; automatic decisions on hub location can be made, topology design is automatic, and the access and National Networking Channels (NNC) are dimensioned.

- **TOD (time of day)**—this is a stand-alone tool based on the NEMO platform, which allows design of global networks. It caters for variable traffic demand across the day in different time zones, thus allowing the designer to take advantage of spare bandwidth in zones which are 'sleeping' by switching traffic from busy 'daylight' zones. TOD can exchange data with the Make NetMaker tool (see earlier).
- **CAT (Cyclone access tool)**—Cyclone is a BT global project whereby points of presence are situated in various countries allowing international traffic to be handled by BT with no intervention by the foreign operator. CAT is another stand-alone NEMO tool which stores global international private leased circuit (IPLC) tariffs and details of the Cyclone points of presence, and has a clustering capability between customer sites and the points of presence which enables the tool to create a topology which best fits the local situation.
- **FBS (flexible bandwidth service)**—this is a NEMO-based tool still under development which will

allow costings of customers' FBS solutions to be obtained quickly and accurately. It will take account of complicated diverse routing rules which affect design and costs.

- **SMDS** (*switched multi-megabit data service*)—Again, a NEMO-based tool which caters for pricing access to the SMDS.

Worldplan

Worldplan is a proprietary tool used increasingly throughout BT for design and costing of voice networks. It has global applications, having access to tariffs for many parts of the world. In terms of user interfaces, it is based on its predecessor Micronap, which could be used only for in-country, rather than inter-country, networks. Worldplan has a 24 hour feature similar to that mentioned for the time-of-day tool previously. This facility allows the user to enter traffic profiles in various parts of the world as local times; the tool then automatically allows for different time zones. Another novel feature is that it can store up to 48 hours of local traffic data; this then ensures that the entire global network can be modelled for at least 24 consecutive hours without needing to compensate for time zone slippage. Traffic profiles can be input manually, or automatically from standard traffic measuring equipment. 'What-if' modelling is another of the main features. Worldplan runs on a PC platform.

NetMod

NetMod is a BT in-house development widely used to design and configure Timeplex multiplexer networks. It can automatically upload and download configuration data directly from the network, and allow a complete sequence of upload, modelling, reconfiguration, and download back to the network to be done without rekeying of data. Although NetMod has performed admirably for several years, the pace of technology and its inability to handle other manufacturers' equipment without major overhaul

has recently led to the decision to withdraw the tool and replace it, probably with NetMaker XA. NetMod has been used successfully to design and maintain networks not only for several major UK customers, but also for BT's own internal data network, the IBTN, run by Global Networks for use by BT worldwide with over 200 nodes in the UK, USA, Far East, Australia, and Europe. NetMod is Sun Sparc-based.

Site audit tool

Several tools in the portfolio are not true design tools, but can assist the planning and design processes. The site audit tool is one of these. The particular application which prompted this development was an inventory management tool for BT units involved in outsourcing communications management, but it could have use in many different areas. The tool is a hand-held computer with a bar-code reader. Before going to site, a service-specific program is downloaded to the tool; this gives menus, options, and parameters which simplify the operator's work on site. For example, if it is known that 90% of the equipment on site is of one of three types, then the options menus can be written to give the three choices as single keystroke selections; if the equipment is the odd 10%, the operator can enter a fuller description. A bar-code label is attached to the equipment, and the code also entered on the tool. At the end of the day, the day's input can be uploaded to a central computer by modem if necessary.

Future visits to site can rapidly add to, or show the validity of the original audit. A manual audit is prone to error because during the significant time it takes to implement, the equipment moves around so that items are missed or double-counted; in the next stage of keying in the paper input, there are also significant errors such as mis-typing. All this is reduced with an audit using the tool because the time-scale is shorter, and there is no rekeying.

Structured requirements capture tool

When a customer approaches BT with a request for a solution, the objective is that the requirements are captured accurately, without the need to return to the customer to ask questions which were not foreseen or forgotten on the first occasion.

Similarly, the requirements capture exercise should not necessarily identify the solution; this should be the result of analysing the requirements and determining which service, technology or mixture of these is best for the individual customer. This might involve modelling the solution on different services and comparing the results in terms of cost, resilience or any of a number of other factors which may be of varying importance to each customer. Since this would mean inputting the requirement parameters to different tools, traditional methods would probably entail a significant amount of rekeying.

The structured requirements capture tool (SRCT) is a front-end customer interface which prompts the system engineer for input and ensures that all the customers' requirements are captured whatever the eventual solution.

SRCT connects directly to tools that design networks, ensuring that all information is captured accurately for all foreseeable scenarios.

Service level design tool

This tool, which is still under development, will be capable of accepting input from the structured requirements capture tool and providing information on the overall direction that might best fit the customer's needs. Often, a customer might want a quick ballpark figure that will enable him to determine whether the requirements which he has just specified can be met with his available budget. This figure will allow him either to commit more expenditure at an early stage, or be less ambitious in his expectations. This has the obvious advantage for BT that practically no resource has been

expended in producing a bid which will be rejected out of hand.

The tool is a case-based reasoning system which holds information on past designs; a search of this data reveals whether the present customer's requirements are similar to historical data. If so, this can give pointers to likely solutions and ballpark costs. As important, if the previous bid was lost or the design floundered for some reason, the same mistakes can be avoided. This particular aspect aligns with one of the BT values of continuous improvement. Updating the database regularly with new cases will ensure that the tool is effectively 'learning'. A key non-technical spin-off of this approach is that it supports a 'partnership' with customers from the earliest stages of negotiations.

ACD modelling tool

Regional Network Planning has design responsibility for the largest networked automatic call distributor (ACD) system in the world, known as *CSCnet*. The network consists of about 60 Northern Telecom Meridian ACDs, many of which provide the primary interface to customers for purposes such as billing queries, fault reports, and sales; however, Northern Telecom's management system allows fully-integrated management of a network consisting of a maximum of about 21 fully-connected ACDs.

To ensure optimum design and strategic utilisation of the network, Network Design Standards and Tools has arranged development of a PC-DOS tool which enables management and design of Meridian ACD networks. The tool incorporates knowledge of the call queuing algorithms used by Meridian ACDs, and allows 'what-if' scenarios to be modelled much more rapidly than the equivalent manual methods, ensuring that configurations and resources can be selected which will provide the best quality of service (in terms of waiting time experienced by a customer) for the installed equipment.

The Future

Network design has traditionally been performed manually, and there is a widely-held belief that it cannot be automated because of its high degree of creativity. Although this may be true to the extent that a totally unskilled operator would be unlikely to create an optimum design using even the most sophisticated tools, the use of automated tools allows skilled designers to design networks in a fraction of the time previously possible. Even more importantly, changes to the original parameters can be input quickly and accurately to give scenarios which can be compared for cost, resilience or grade of service.

This article has briefly introduced one of the fastest-growing areas of BT's global business, and has described some of the ways in which the Network Design Standards and Tools unit is addressing the challenges of customer network design while ensuring that the limited pool of skilled resources is not wasted on mundane tasks. In order to be competitive, the use of automation to implement network design processes must increase. At the same time, the limited resource pool must not be allowed to erode further. This can only be achieved by ensuring that competent network designers are recognised, registered, and rewarded.

Acknowledgement

Acknowledgement is given to Make Systems (UK) Ltd for their permission to use the figures relating to NetMaker XA.

Biography



Paul Warren
BT Worldwide
Networks

Paul Warren joined the Post Office as an apprentice in 1966. After graduating from Loughborough University of Technology in 1974, he worked in several areas of PO and BT Headquarters, including analysis of fault statistics on TXE4 exchanges, Monarch PABX development, and testing and validation of computer systems. More recently, he has been involved in the evolution and strategy of internal networks and now works in Global Networks, where he is responsible for the identification and provision of automated tools for use in network design, especially in the global customer arena.

The Design and Implementation of the BOXER Network

The BOXER network, one of the largest private networks in the UK, provides nationwide connectivity to RAF and other military and government users. This article describes various aspects of the network, which has been designed and implemented by Syntegra, BT's systems integration business against exacting Ministry of Defence requirements.

Introduction

The BOXER network has been designed and implemented by a dedicated project team within Syntegra, BT's systems integration business, against a Ministry of Defence (MoD) contract placed in 1987, with a scheduled completion and final handover date to the RAF of June 1996.

The BOXER network provides full nationwide connectivity to RAF and other military and government users as part of one of the largest private networks in the UK. The network forms the transmission element of the RAF fixed telecommunications system (RAFFTS). As a military network the system has been designed to meet the specified requirements of security and survivability.

BOXER consists of over 200 individual sites, the majority of which are further broken down into sub-sites. The basic topology is shown in Figure 1. As a standalone MoD network, one of the primary design aims upon BT was diversity from existing BT sites wherever possible.

BOXER uses both optical-fibre and microwave radio systems to provide the bearer network, and operates at hierarchical data rates up to 140 Mbit/s. A general outline of the network is shown in Figure 2. The network consists of over 4.5M channel kilometres of 64 kbit/s digital paths.

One of the first decisions made was how the project should be organised, given the size, complexity and geographical spread of the network. It was concluded that the use of a phased approach, with each

phase consisting of six-month periods, provided the most manageable solution. Using this as an objective, the project was broken down into specific design/implementation elements, the functions and operations of each being as follows.

Network Design

The project team recruited network radio and cable planners from what is now BT Worldwide Networks to carry out the detailed design work.

The design of the microwave network used the experience built up by BT Research at Martlesham Heath for the main BT network, and relied on the proven computer-aided design packages developed there. Among the software packages used heavily is the Peacemaker path profiling and interference study suite of programs. Having completed desktop studies using these systems, detailed route and site surveys are undertaken in order to confirm the equipment configurations to be supplied, tower availability and antenna mounting access at the required heights.

Cable planning operates in two major areas, on-site and access. The on-site cable planning identifies the routes required between the microwave radio access site and that of the traffic end-user (see Figure 1).

Access planning concerns those routes where fibre provides the transmission medium into the site via dedicated duct installed in public highways in a manner identical to BT and other utilities. As such, the requirements of the New Road and




Figure 1 – BOXER topology

Street Works Act and its predecessor the Public Utilities Street Works Act (PUSWA) apply and need to be considered and implemented in addition to project specific requirements.

The culmination of the above planning activities is the identification of all elements required to be provided on a site in terms of civil works etc. in order for the project to be implemented. This information is embodied in a site-specific civil-works requirement document, which is issued to the MoD or the site owners by the project team with sufficient lead-time to ensure that works are completed by the planned site implementation date.

The final planning activity concerns traffic and network sizing: the overall network traffic requirements and connectivity are the responsibility of the RAFFTS planning team; the physical implementation in terms of capacity and multiplexing lies with the BOXER project team. The results of this planning effort are specific multiplex and equipment plans for each site in the network; these are used to define the equipment fit for each location (see later).

Equipment

The BOXER network utilises over 30 major items of equipment ranging from power supplies including stand-by generators and rectifiers, transmission equipment (multiplexes through to microwave radio systems) and equipment housing. Of these the most

important is the *transportable prefabricated building (TPB)*. This is the cornerstone of the project as it provides security and survivability protection for the transmission and ancillary items of equipment. The TPB in its final configuration provides environmentally-controlled power and

equipment mounting areas separated by a fire wall, each with its own means of external access. External protection for the TPB is provided by insulated moulded weatherproof GRP panels, which are designed to blend with the location by means of various finishes; for example, brick or plain coloured as

Figure 2 – Boxer network national coverage



standard, or others to meet special requirements.

The overall dimensions of the TPB (7 m × 3 m × 3 m) allow for transportation by road, and the design is provided with suitable lifting points not only for craning on and off the low loaders for road delivery, but also for helicopter delivery to remote locations.

The procurement of all items of equipment is based on competitive tendering exercises managed by the BT Group Procurement Services personnel allocated to the project. Each item of equipment required is based on standard commercially available items wherever possible to minimise specific development programmes and products. These are then evaluated against project specific facility requirements contained in BOXER specifications. These documents, allied to specifications defining the BOXER project requirements in both the commercial and quality-assurance areas, form the basis for the tender adjudication exercises.

The procurement exercises operate under a strictly formalised methodology based on achieving an overall value-for-money solution covering whole-life costs. These exercises resulted in over 500 individual tender responses to the various requests, and, after adjudication and acceptance by the MoD, resulted in contracts worldwide with various major telecommunication suppliers.

Implementation

The size of the network and the time-scales required necessitated a structured programme of implementation. This programme consists of three distinct elements:

Element 1 (Release Point 1)

This element concerns the procurement and delivery of the individual equipment items from the manufacturers.

Each equipment type used within the project is subjected to a specific type-approval exercise to confirm compliance with the network performance requirements. Based on these

tests, a set of release criteria are then agreed with the companies; these criteria are used by the project's quality assurance (QA) personnel as a basis for acceptance of each batch of equipment. Based on their performance in this area, and to maximise effective use of the project QA resource, companies are encouraged to achieve 'delegated release status' with batches of equipment being released under certificates of conformance issued by their own internal QA organisation.

Delivery of the procured items is directed to one of two specific areas: either the centralised installation facility (Element 2 below) or direct to site in the case of items such as antennas and associated mounting steelwork.

Element 2 (Release Point 2)

This element concerns the site-specific configuring of the TPBs. In order to achieve this, a centralised installation facility was developed and operated by Fujitsu Fulcrum Communications Ltd. under a BT contract. This facility is required for two reasons:

- The range of items to be installed in the TPB with the associated fittings and fixings precludes this as an on-site operation.
- The installation rate required (an average of 44 TPBs per year) demanded significant manpower with limited scope for efficient utilisation given the wide geographic spread of sites.

The manufacturing cycle through the centralised installation facility follows a standard sequence, which is now described.

The configuration in terms of equipment fit is identified as part of the network design and planning functions of the main BOXER team. This information is then transmitted to Fujitsu Fulcrum. Fujitsu Fulcrum is responsible for designing and fitting out the TPB under the overall management control and supervision of the project.

Initial fitting out includes the provision of overhead ironwork and trunking to support the TEP-I(E), CEPT Slimline or other equipment housing requirements. AC electrical power wiring is then installed to complete the basic fitting of the building. The building is now ready for the equipment to be installed. Radios utilising CEPT Slimline practice and the TEP-I(E) racks containing the multiplexes and optical line terminating equipment are fixed to the overhead ironwork.

The power equipment consists of two primary elements, both based on standard BT products. The first, installed in the power room of the TPB, is the stand-by generator (PS 4006 equivalent) and its associated ventilation/exhaust system, and the second is the rectifier and stand-by battery system EP2008B installed in the equipment room. Once the physical installation of the equipment is completed, the remaining power and transmission cabling can be undertaken. When this is completed, the configuration and connectivity of the equipment elements via digital distribution frame (DDF) cabling is matched to the defined traffic connectivity requirements for the site the TPB is allocated to.

The final element of the Fujitsu Fulcrum manufacturing cycle is testing (see later). The design of the Fujitsu Fulcrum facility allows for the benefits of a full factory manufacturing cycle to be implemented. This cycle takes 12 weeks to complete for each TPB, and at the production rate required for the project, 12 TPBs will be in manufacture at any one time.

This structured process allows for maximum throughput in a controlled manner, and provides for efficient utilisation of the available manpower resources and skills.

Element 3 (Release Point 3)

The final implementation element is installation on site and associated testing/commissioning. BOXER sites are located throughout the UK mainland. Although road access is

Figure 3—Typical transportable prefabricated building installation

achievable at all sites, in several cases four-wheel drive or special-purpose vehicles have to be used. In these cases, delivery of the TPB on its standard road low-loader transport can prove impossible; a helicopter or other specialist delivery system is then necessary. A fully equipped TPB weighs between 9 and 10 tons, which is just within the maximum load of an RAF Chinook depending on how far the building has to be carried from the access to the site. As stated earlier, the TPB is designed with appropriate lifting points to facilitate this delivery method, which has already been successfully utilised, and is scheduled for a number of other site deliveries before the project is completed.

Each site must be equipped with appropriate mounting facilities for the TPBs. These foundations and associated other civil works, for example, ducts etc., have to be completed before deliveries are commenced. The off-loading of the TPB, as is the mounting of antennas on microwave towers, is managed by BT Worldwide Networks via a service level agreement. Their responsibilities cover the off-loading and safe bolting down of the TPBs, and the installation, panning and testing of the antenna systems.

Once the TPB is bolted down, the responsibility for installation reverts back to project personnel. The building is unpacked and the task of recommissioning starts. This occupies many functions: connection of the TPB to its primary 240 V AC power supply, the connection of the stand-by generator to its fuel supply, fitting of weatherproof cowls over the environmental ventilation grills, and reinstallation of the transmission equipment modules/cards into their rack positions. These items are demounted and packed for transportation purposes by Fujitsu Fulcrum before dispatch as a general policy, although tests proved that transportation of a fully equipped building is a viable option for most of the UK sites.

Once the physical reinstatement of the TPB is complete, power can be connected to the equipment to allow



testing to take place. The objective is for the primary power to be connected to the TPB such that the environmental control systems (heaters/fans) are operational on the day of delivery; a stable environment for the implementation personnel will therefore exist during testing and commissioning.

Once all the testing is completed, the building is cleaned and the internal cabling connections on the DDF are returned to the traffic design requirements before handover to the customer.

Figure 3 shows a typical TPB installation.

Performance

The performance of the network is assessed in two ways, firstly in terms of individual equipment elements and secondly in terms of the network itself.

Equipment testing

The testing strategy adopted in the project follows similar lines to that for implementation: the concept of reducing structured testing through the three stages identified as Release Points 1, 2 and 3 allows for minimum staffing levels and time on site.

Release Point 1

This testing covers the factory release of the individual items from the manufacturers. The tests are derived from those used for design acceptance, and are used by the project QA personnel to assess each company's performance against the company's quality plans. Formal acceptance of batches of equipment requires either on-site visits or

acceptance of certificates of conformity issued by the individual company's QA authorities. For the microwave radio systems, these tests have to be carried out on equipment configured as 'hops' and incorporate both equipment and hop testing. This is made possible since equipment is ordered on this basis, and not on individual terminal ends.

Under the contracts the manufacturers must hold full records of all the test results, as reference sources for both BT and the MoD quality assurance departments, and are used as reference data for the regular audits identified in the agreed quality plans.

Release Point 2

Two stages of testing take place at the centralised installation facility: goods inwards and goods outwards or final release.

- **Goods Inwards Testing** Owing to the complexity and numbers of different items involved within the project, the need for initial testing upon delivery to the CIF to check for damage in transit and basic functionality is essential. On many of the items, for example, multiplexes and line terminating equipment, this utilises standard test rigs; however, for the larger more-complex items, particularly the radios, full assembly becomes a prerequisite of testing. Because of this, Fujitsu Fulcrum modified the goods-inwards testing strategy for the radios to incorporate hop testing as a more effective method of proving performance. Certain items of equipment can be functionally tested only after installation in

the TPB, particularly the COMPAC stand-by generator. Goods-inwards testing of this and similar items is therefore limited to signs of visible damage. Having some degree of certainty that the individual items are functional, installation in the TPB then takes place.

- Goods Outwards Testing** This testing is undertaken once full installation of the equipment elements is completed. As identified earlier in this article, each TPB is configured on a site-specific basis equating to the traffic levels both through and terminal. Thus, by means of the DDF, equivalent traffic paths can be constructed within the building to verify performance and interoperability between the equipment items. Looping at IF on the radios allows for at least partial verification of those items allied to single terminal testing of parameters such as output frequency, power and spectrum. With the generator installed in the TPB, along with its associated input and exhaust systems, verification that it will start and run satisfactorily can now take place. Also the wiring of the TPB is tested and certified as being in accordance with the IEE wiring regulations.

Once all this testing is completed, a certificate of conformance is issued and the TPB is prepared for dispatch to site. As with Release Point 1 testing, all test results are held by Fujitsu Fulcrum to be used as a basis for ongoing project and QA audits.

For those items of equipment delivered direct to site, for example, waveguide and antenna systems, optical-fibre cables etc., the requirements for Release Point 2 testing are waived.

Release Point 3

On-site testing is limited to that necessary to confirm that no damage has occurred in transit for those items delivered from the centralised installation facility, to verify the performance

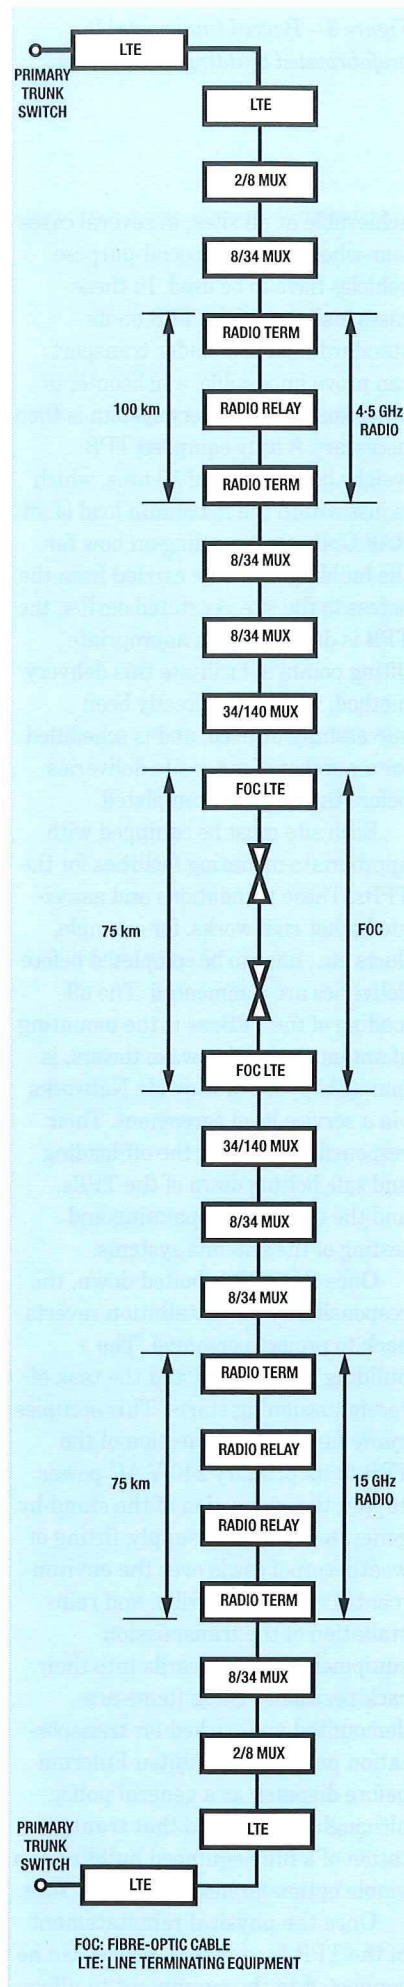
Figure 4—BOXER hypothetical reference digital link

for those items delivered direct to site, for example, antenna and feeder systems or installed optical-fibre cables, and to test overall radio hop/optical-fibre sections. BT Worldwide Networks install and commission the antenna and waveguide feeder systems and pan the links. Gas-barrier-to-gas-barrier loss and swept-frequency return losses provide acceptance criteria for this work. Testing of the equipment within the TPB or end-user cabinets is carried out by the project implementation managers and engineers against a set of procedures defined and agreed with the MoD. Hop/fibre section testing is carried out against targets derived from the network performance objectives described in the next section, and these results form the basis for final acceptance by the MoD.

Network performance

The BOXER network performance is assessed by using the principles defined in CCITT Recommendation G.821. With the mix of data rates between 2 and 140 Mbit/s, and the configuration of the network, sections of the 'high, medium and local grade' elements of the hypothetical digital reference link (HRDL) defined in G.821 need to be considered. By using G.821 as a basis, a specific BOXER HRDL (see Figure 4) has been derived and agreed with the MoD and their supporting technical experts. From this, performance targets in terms of errored seconds and degraded minutes are derived on a per-kilometre route basis. The performance targets for each link within the network can then be calculated relative to its path length.

The network is considered as a series of testable sections (see Figure 1) and performance is verified on each of these. Since each section's performance is calculated on a per-kilometre basis, chaining series of sections together to form logical links allows the overall end-to-end section objectives to be met. The acceptance of each section by the MoD as they are completed is based on the results



in annual assessments, BOXER has comfortably exceeded all its targets by a considerable margin

of a final 200 hour stability run against calculated targets allied to confirmation of testing at Release Points 1 to 3 as defined above.

In terms of the BOXER network, overall performance lies between the 'high and medium grade' CCITT objectives; however, by planning the network solely against the 'high grade' criteria, the customer gains benefit in terms of reliability and performance.

The performance of the network against the objectives is the subject of continuous monitoring throughout the implementation programme by the BOXER team, and the RAF provides results on the transferred sections. So far, in annual assessments, BOXER has comfortably exceeded all its targets by a considerable margin.

Long-Term Support

Having successfully completed the implementation of a site and transferred it to authority (MoD) control, the BOXER project team provides a range of support services to the RAF via the MoD project office.

These services cover a range of requirements from site-specific documentation through to spares provision and, to assist the RAF maintenance personnel, a 'help desk' facility is now operating to allow them access to the expertise existing within the project team.

As more and more of the network is transferred to authority control, the need for and growth of these support services is assured.

The support area of the project team has one other key role to play. As with all major networks, change is a fact of life and BOXER is no exception to this rule. The support team provide the focus for all requests from the authority for studies to assess the impacts of various changing network requirements. These studies are costed, and, once funding is agreed, resources from both within the project team when available or from other Syntegra units are allocated to the work. The outcome of most of these studies is revision of the overall

implementation contract and this work provides an important source of additional revenue.

Management Structure

The basic project structure consists of a central project team containing the management and planning functions and separately located implementation teams, which were originally located in Bristol, Manchester and Inverness. As part of a recent rationalisation exercise, the opportunity was taken to concentrate on two BT customer system sites, the central project team at Fleet, and the national implementation team in Leeds.

The central project team has been set up to operate so that the security requirements of the project can be adequately addressed. Thus a team containing members with all required skills and disciplines has developed; these skills are not only in the technical area but cover procurement, contracts, quality assurance, commercial and support.

The size and complexity of the project has necessitated development of structured reporting procedures and communication channels between the various elements. Formal communication channels and structures between the MoD and the project, the project and the manufacturers and other agencies, and between the various project offices have been put in place and provide the reference source for all decisions made. Strict configuration control procedures under a dedicated manager allow for full confidence in the system design, project documentation and database accuracy and tracking. The definition of agreed management reporting structures allows for both internal and customer facing reports in all aspects of the project planning, implementation, milestone achievement and finances to be generated and presented on the specified dates.

Conclusion

The BOXER project has provided Syntegra with an opportunity to

demonstrate its abilities in the management and implementation of a complex multidisciplinary high-visibility high-impact project. The skills and procedures developed have allowed and continue to allow the project to meet its major objectives within time and budgetary limits. They allow for the variations which occur to the planned programme to be accommodated in the best manner for both BT and the customer and with the minimum impact wherever possible.

These skills, while they have been developed within BOXER, are not specific to it; many other BOXER-type projects still remain to be won and the skills are equally appropriate to other types of major project and are recognised by Syntegra and included within its skills management systems and processes.

The completion of the implementation of the project is scheduled for June 1996, and the project team is fully committed to and have every confidence in meeting the date.

Biography



Alan Pritchard
Syntegra

Alan Pritchard is Project Manager Design for the BOXER project, having joined it as technical

manager in 1988. He has worked in various aspects of the transmission development department of BT Headquarters including management systems, standards where he represented BT on CEPT study groups and other standards bodies, before finishing in the Microwave System Group responsible for the introduction of digital systems in the 2, 4 and 19 GHz bands as well as acting as a consultant on numerous Telconsult studies. He is currently managing the development, certification and implementation of a network management system for the BOXER project.

The 'Value' of the Telecommunications Engineer in Eastern and Western Europe

Throughout the western world, both in the USA and Europe, the telecommunications industry has undergone dramatic change arising from political, technological and social-economical trends. The outlook for the future is continuous change driven by global customer needs; a paradigm shift that will not be a panacea for all. Rapid transition from PTT to telco will be made by companies that have telecommunications engineers that are not only good engineers, but managers and leaders of the highest calibre. The 'value' of the modern telecommunications engineer will be greatest for those that understand the complex interactions between three key elements: technology, the market, and commerce.

Market Trend

Within the western world, the trend is for telecommunications operators to move from PTT to telco and if they are to remain as a key player then they will have to become a 'super telco'. These transitions are essentially driven by customer demand, increased competition and regulation. While there is a general reluctance by some PTTs to accept full liberalisation and possibly privatisation, it is generally accepted that all will have to be managed as competing telecommunications businesses—a 'telco'. Survival will depend on them adopting world-class techniques in revenue generation, while taking mature and bold steps to reduce costs—the 'super telco'.

The competition directorate (DG IV), of the European Commission, has begun to show an increased tendency to intervene in forcing the rate of change. If the drive does not come from Brussels, then it will be from the multinational companies who will become increasingly dependent on modern telecommunications infrastructures. The prize for the serious players is a slice of the European Telecoms Market that is estimated to be worth ECU 110 billion¹.

This article is based on a paper presented by the author to the 1994 FITCE Congress, in Dresden. David Corrie, a Member of the IBTE, won the Best Paper award for his paper. The theme of the Congress was 'The European Challenge to Telecommunications in the East and West'.




Figure 1—Development from PTT to telco

Clearly it is not a case of **if** the PTTs become 'telcos', but **when** (see Figure 1)!

The rate of development of the PTTs will be influenced by such factors as political, technological, regulatory, and social-economical trends. It is not important to identify or even speculate on which graph applies to which PTT. It is more important for the telecommunications engineer to appreciate the range and identify the one that applies to his or her PTT. Only major events will change the rate of development. An example of this would be the unification of Germany, which, it could be argued, was a major factor in influencing DBP to opt for privatisation.

With regard to Eastern Europe, the PTTs have a similar start point in that they are government-run administrations. Longer term, their need to encourage investment and hunger for 'know how' could position them ahead of many other PTTs in




Figure 2—Development of the Eastern European PTT

central Europe. Rather than a gradual move towards 'telco' status, there could in fact be a step increase in the way they manage their telecommunications business (see Figure 2).

The Role of a Telecommunications Engineer within a PTT

The 'value' that an organisation places on the telecommunications engineer will be directly related to the role. To be of highest value the role should be seen as enhancing the purpose of the organisation.

Within the PTT, the telecommunications engineer has a technical role, or in the extreme could become an administrator of a system that is run on simple business lines.

A simplified business model of the PTT is shown in Figure 3. In the extreme, the PTT operates as a bank for the government.

The overall internal structure and environment tended to condition and constrain the engineer. The absence of flexibility tended to suppress

Figure 3—Simplified PTT business model

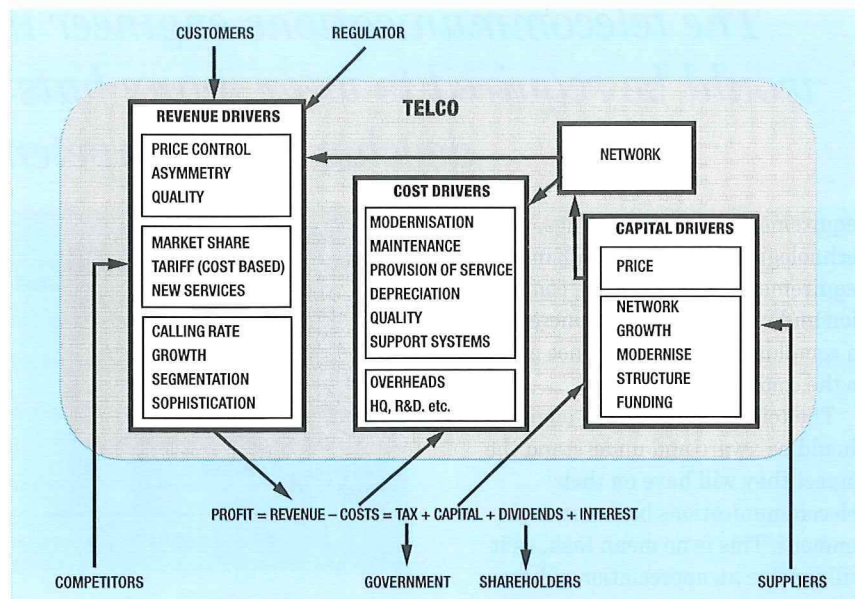



Figure 4—Simplified telco business model

initiative and enterprise. Limits on authority contributed to delay in decision making as did the degree of centralisation. Overall effort tended towards sustaining the system and fulfilling ministerial and parliamentary needs rather than those of the subscriber².

Where the engineer had control, it was generally exercised in the physical technologies that were to be introduced into the network. The engineer could differentiate between supplier equipment types and determine

little resemblance to the PTT business model.

In its simplest form, it consists of cost and revenue drivers, but it is dominated by the network. For the 'telco', the market is the most important external driver, because it generates revenue. However, market share is influenced by competition, which itself is manipulated by the regulator, and influenced by government policy and the economy. Market share is also driven by tariffs, quality and the products and service portfolio.

The emergence of the 'telco' has resulted in a 'value' shift for the traditional telecommunications engineer

suitability for network compatibility and CCITT conformity. However, the engineer's recommendations could be unacceptable in terms of other non-technical criteria. Essentially, the role would be well defined and stable.

The Role of a Telecommunications Engineer within a 'Telco'

The emergence of the 'telco' as a telecommunications business with government and/or private shareholders has resulted in a 'value' shift for the traditional telecommunications engineer.

The 'telco' business model (see Figure 4) is a multifaceted complex and dynamic structure, which has

lio, all of which are heavily influenced by the network.

An example of this complexity is illustrated in the tariff for a new service. This is determined by the cost of providing it, the anticipated revenue to give the required return and any regulatory restrictions. The cost is dominated by the network drivers, one of which is the price of equipment, which is influenced by supplier competition and rationalisation. It is therefore evident that the drivers are independent and interactive. In its simplest form, it is the complex interaction between the market, technology and commerce. In addition to these complex interactions, there is a further dimension—that of time. Customers'

The telecommunications engineer in the 'super telco' would be required to wear many hats and be prepared to develop new competences

requirements, service offerings, technological innovations, business requirements, globalisation, competition and regulation will all increase in complexity with time, hence adding to the complexity issue.

The telecommunications engineer should be aware and understand the impact they will have on their telecommunications business environment. This is no mean task, as it will involve an appreciation of how these parameters interact and drive the telecommunications business. Essentially, to be of highest value to the 'telco' the telecommunications engineer will have to have more than just a good understanding of the technological issues.

The new criteria for the telco organisational effectiveness will change to flexibility and adaptability. Job design will stress personal growth and responsibility. Decision making, control and goal-setting processes are decentralised and are shared at all levels of the organisation.

Essentially, the telco is a business and as such must be managed rather than administered³ (see Figure 5).

As the 'telco' gathers momentum in the complex and highly competitive world of telecommunications, it moves into the final stage of metamorphosis—the 'super telco'. In this organisation, which has the sole aim of becoming best of breed, the streamlined workforce will be required to be leaders of the highest calibre.

For the 'telco' to survive it has to have a workforce who not only

Figure 5—The rise of the engineering manager





Figure 6—The emergence of the engineering leader

understand the telecommunications industry, but have to be part of company culture that encourages many new ways of working. To be of highest value telecommunications engineers have to be able to manage ambiguity in an age of uncertainty; to sell as a way of life; to be comfortable with risk taking; and, overall to be seen as leaders who do the **right** things rather than just doing things right (see Figure 6).

Role definition is broad and subject to change. Career progression may be limited in what will be a downsizing organisation, and development may involve working in different parts of the telco (non-engineering). Promotion would be based on merit as determined by performance appraisal³.

The telecommunications engineer in the 'super telco' would be required to wear many hats and be prepared to develop new competences.


The 'Value' of the European Telecommunications Engineer

It is a reasonable assumption that a buyer (the PTT, 'telco' or 'super telco') will choose the supplier (the telecommunications engineer) that offers greatest added value to the buyer's organisation. This value will be influenced by the market, in which the buyer may be: a sole trader (the PTT) or one of a number of traders (the 'telco'). In essence, for the telecommunications engineer to be of

value to the telecommunications operator, he or she will have to adapt at a rate that is in-line with the changes that are taking place within the organisation: engineer to manager and eventually to leader.

The challenge for the telecommunications engineer is the ability to recognise where his or her organisation is in the continuum from PTT to 'telco' to 'super telco', and then matching the personal skill set to the needs of the business (see Figure 7).

Figure 7—Matching skills to the needs of the business



In the extreme, the 'super telco' could be looking for telecommunications engineers who have some, if not all, of the following competencies:

- Outstanding customer focus
- Good commercial and business awareness
- Good performance and results
- Leadership, motivation and people management
- Good teamworking skills
- Effective communications and impact
- Financial awareness
- Global awareness
- Strategic vision and direction setting
- Creative thinking and innovation
- Good planning and organisation/project management
- Good problem solving and decision making
- Self management and personal development

Equally, it would not pay to be too proactive and become a leader before the organisation is ready; the behaviour may be seen as being radical and disruptive.

The telecommunications engineer from Eastern Europe is beset with a similar, but different, challenge: the PTTs are based on the same business model as Western European PTTs, but they are making step-function leaps toward becoming a telco. This has been brought about by massive investment, in the telecommunications infrastructure, by North American telcos.

The transition from engineer to manager, and eventually to leader, may be so rapid that it could leave the Eastern Europe telecommunications engineer behind. Individuals could be left wondering what they have to do to be of value to the organisation. On the positive side, with the right training and development, the Eastern Europe telecommunications engineers could find themselves a much sought after commodity in the West!

Conclusion

The traditional European telecommunications engineer is about to embark on a career of continuous change, driven by competition. For those that have made the transition to the 'super telco' environment, they will find it littered with many double-edged swords. For in their Utopia they will discover a financial freedom of a size they could have never envisaged from their days in the PTT. But will the telecommunications engineer lack the knowledge needed to invest it in a way that helps the growth of the 'telco'?

In essence, for telecommunications engineers to be of the highest value to their organisations, they will have to understand where the business is in its growth from PTT, 'telco', to 'super telco'. They will need to look for personal training solutions that map onto this change. Care must be taken not to advance too rapidly, as the

resulting behaviour may not be tolerated.

In the case of the Eastern European telecommunications engineers, they need to watch out for step changes in the growth of their organisation, and look to the infrastructure suppliers to provide the integrated training solutions.

Finally, the engineer has, and will continue to have, the highest value in the PTT, but it will be the leaders that will reign supreme in the 'super telco'.

References

- 1 Telecoms Markets, Financial Times Business Enterprises (FINTEC).
- 2 STEWART, ROSEMARY. The Reality of Management. Pan Business Management, 1979.
- 3 LYSAGHT, J. Facing up to 1992: FITCE and the Engineering Profession. Telecom Eireann, 1989.

Biography



David Corrie
Head of the BT
Telecommunications
Master's Programme

David Corrie is Head of the Telecommunications Master's Programme (formerly the Telecommunications Staff College). He is presently working on the next-generation programme, based on BT's multimedia products, with the aim of developing BT's leaders around the globe. After graduating in 1979, David took up his appointment with the System X launch team. Other key assignments include information technology management, undergraduate development programmes, a major joint venture training programme, and international training support. David is a member of the Institution of Electrical Engineers, the Federation of Telecommunications Engineers of the European Community (FITCE), and a Fellow of the Institute of Personnel and Development.

33rd European Telecommunications Congress

*This article reports on the
33rd European
Telecommunications
Congress, held in
Dresden, Germany, from
29 August–3 September
1994, organised by the
Federation of
Telecommunications
Engineers of the
European Community.*

Introduction

The 33rd annual congress of the Federation of Telecommunications Engineers of the European Community (FITCE) was held in Dresden, Germany, between 29 August and 3 September 1994. Dresden, with its recent history, proved an appropriate place to host a congress with the theme 'The European Challenge to Telecommunications in East and West'.

Once again the congress proved to be extremely popular with European telecommunications engineers. It was attended by over 600 (participants and their partners) from the European Union, Switzerland, Poland, the Czech republic, Slovakia, Austria, Finland, and Rumania. Over 17 countries were represented by professional telecommunications engineers, with Germany taking the lead in sending some 120 delegates to the congress. The UK was repre-

sented by a small group consisting of representatives from BT, GPT, and AT&T.

Some 44 papers were chosen by the selection committee, headed by Dipl.-Ing. Hermann Gabler for presentation at the congress. The awards for Most Promising Young Engineer and Best Technical Paper were both won by BT people.

Opening Ceremony

The congress began in true FITCE cultural style with Mozart's flute Quartet in D Major played by the Dresden Philharmonic Ensemble. Dr. Wolfgang Bötsch, Secretary of State of the Federal Ministry of Posts and Telecommunications, formally opened the congress. In his warm word of welcome he stressed that Dresden was an excellent choice of venue, which was not only the capital of the Saxon Region, but also the centre of rapid development of

Some of the Congress delegates at the technical conference, Dresden



telecommunications in Eastern Europe. He added that there was a need to continue the effort to improve the supply of telecommunications service which would create the right environment to accelerate the economic upturn. The Deutsche Bundespost Telekom would be driving forward with liberalisation and privatisation on the path to harmonisation of the European States.

Herr Kurt Biedenkopf, President of the free states of Saxony, said that the theme of the congress was particularly appropriate for the new federal states. He commented on the advances that had been made in the network, with two-motion selectors being used up to 1991. He said that trade follows the telephone and this had certainly been true in Dresden. He had seen an upturn in the economic growth of the region. He concluded by saying that money alone would not have been enough in developing the telecommunications infrastructure, great emphasis was placed on the transfer of 'know-how' to the new states.

Herr Helmut Rieke, Chairman of the Deutsche Telekom Management Board, stressed the importance of the congress, of the location and, in particular, the unification of telecommunications in the East and West of Germany. His speech focused on putting the customer first and the need for strategic alliance in DBP's strategy for international growth. Herr Rieke stated that there was an increased need to get the balance right between collaboration and competition, and that the congress was a good example of collaboration between European operators.

Other speakers at the opening ceremony were the Lord Mayor of Dresden, Dr Herbert Wagner, and the President of FITCE, Prof. Dimitri Kouremenos.

Technical Papers

The 44 formal papers were presented in the following sessions:



- Strategies and Policies;
- New Developments;
- Projects and Experience; and
- East and West Cooperation.

The Strategies and Policies sessions of the congress were chaired by France's Antoine Jousset. Some nine papers were presented on topics as diverse as meeting the needs of European multinational customers, standards, regulation, superhighway, manufacturing satellite communications, personal networking and a futuristic look at telecommunications.

The honour of presenting the first paper went to the United Kingdom, with Peter Cogley from AT&T who reminded the audience of the need to serve the multinational customer. He argued that if the multinational customers are satisfied then all sectors of the European economy will benefit, but this commercial imperative must be met by the European telecommunications industry providing a suitable infrastructure to support this development. The general thrust of all papers in this section was that of growth of telecommunications in Eastern and Western Europe, and factors that need to be considered. This theme was emphasised by Dr. Theodor Irmer of the ITU, who stressed the importance of global standardisation and its impact on the development on Eastern and Western Europe. Overall, it was a very good opening section to the congress which was complimented by a high standard of supporting papers.

The New Developments session was chaired by Dr Augusto Vighi, of

Italy. This session had an interesting mix of papers that touched on private mobile networks; multimedia, as a method of improving cooperation between the East and West; a provocative paper on the 'value' of the telecommunications engineer; an overview of a 2.5 Gbit/s synchronous digital hierarchy (SDH) system used in Italy, and a brief introduction into the working of ETIS (European Telecommunications Informatics Services). There was no common thread throughout this session, with all papers designed to either update or stimulate further debate. A paper that typified this latter category was that of Bavco, Teunissen, and Bauerfeld, and their view on how multimedia teleservices could be used to provide advanced development for Eastern European countries.

The versatile Bill Medcraft, from BT Telconsult chaired the Project Experience session of the congress. His versatility was ably demonstrated when he had to chair his own presentation! This section brought with it a flurry of project managers who provided the audience with an excellent insight into the management programmes that have been undertaken in the East. All speakers commented on the achievements, lessons learnt and frustration of doing business in Eastern Europe. The projects were of varying sizes, none of which compared to that described by Reiner Kostka, DBP Telekom, Germany: the rebuilding of the telecommunications network of East Germany. He outlined how DBP Telekom took one of the oldest networks in the world and transformed it into the most advanced in some five years. He concluded his

presentation by saying that all local networks would have ISDN access by 1995; all local exchanges would be digital by 1997; and they would continue to provide passive optical access network to residential customers, in selected areas.

After a day of technical visits, the conference theme turned to East and West Cooperation, under the chairmanship of Dipl.-Ing. Hermann Gabler, Germany. Essentially, this was an opportunity for the telecommunications manufacturers and consultants to provide an overview of their experiences in supporting the East to rebuild their telecommunications infrastructures. An excellent insight into the 'troubles and difficulties' when consulting in the East, was given by a 'game keeper turned poacher'—Stefan Kolev. Stefan is a Bulgarian, who had been on the receiving end of many East/West support projects, and he now works for a consultancy firm in the Netherlands.

Overall the standard of the papers and presentations continue to improve, which is due partly to the selection committee and the hard work of the presenters. The congress proceedings is a well produced and informative document.

Workshops

For the first time at a FITCE congress, workshops were run as parallel sessions to the main technical conference. Three workshops were held with the following themes:

- transfer of know-how between Eastern and Western Europe,
- mobile radio, and
- networking.

Presentations in the Know-How Transfer workshop ranged from the foundation of a technical university at Poznan, and cooperation between Deutsche Telekom's technical institute and corresponding institutions in

Central and Eastern European countries, to the improvement of telecommunications infrastructure in the eastern parts of Germany, and the establishment of a modern telecommunications infrastructure in Kazakhstan.

In the Mobile Radio workshop, presenters demonstrated the advantages of mobile radio in setting up a network with speed at a moderate cost. Special mention was made about the competitive edge that GSM networks have over existing analogue networks because of their uniform worldwide standards and regulations and allowed international movement without adversely affecting the national operator's responsibility. Various difficulties encountered by the speakers while working in the Eastern European countries were also mentioned.

Speakers in the Networking workshop described their involvements in network modernisation.

The experiences gained from the workshops can be summarised. The operation of a network and the provision of service do not always require the latest technologies. In most cases, the market has first to be developed before any realistic demand estimate for a new service can be given. For customers, it is advantageous to have competition among the network operators providing the same type of service. From the viewpoint of an operator or a (super) telco, however, more than two operators are not at all advantageous; they hamper the development of the telecommunications network. This may also be the case with too many equipment suppliers. Additional costs of basic and advanced training, as well as for additional interfaces, could use up the price reductions obtainable from the suppliers.

Congress Awards

This year, the FITCE Comité de Direction awarded two prizes for congress papers. The prize for the Most Promising Young Engineer went

to Jon Chalmers (BT, Development and Procurement (D&P)); and the prize for the Best Technical Paper to David Corrie (BT, Development and Training).

Jon Chalmers' paper 'East Meets West in the Infosphere' argued that telecommunications has the potential to break down national barriers on a scale that has never been seen before and that the cultural and commercial implications of such a change were enormous. The paper, which was jointly written by Jeremy Barnes (BT, D&P), Ian Pearson (BT, D&P) and Jon Chalmers (BT, D&P), examines globalisation of supply and demand for knowledge workers; new networks and how they may give the East economic supremacy; new cultural synergies; and the development of global 'spaces'.

David Corrie's paper 'The "Value" of the Telecommunications Engineer in Eastern and Western Europe' examines the key 'drivers' that are changing the role of the telecommunications engineer. David argued that the 'value' of the modern telecommunications engineer will be greatest for those that understand the complex interactions between three key elements: technology, the market, and commerce. Throughout the Western world, both in the USA and in Europe, dramatic change has occurred within the telecommunications industry that has arisen from political, technological, regulatory, and social-economical trends. The outlook for the future is continuous change driven by global customer needs; a paradigm shift that will not be a panacea for all. He concludes that rapid transition from PTT to 'telco' to the 'super telco', will be made by companies that have telecommunications engineers that are not only good engineers, but managers and leaders of the highest calibre. (An article based on this paper is included in this issue of the *Journal*.)

Technical Visits

An integral part of the congress is one day of technical visits. This year's

visits took delegates to one of the following companies: Alcatel/SEL (fabrication of printed circuit boards and power supplies), Philips Kommunikations Industrie AG (manufacture of digital transmission equipment), Siemens AG (manufacturing of transmission and switching equipment), ANT Nachrichtentechnik Radeberg GmbH (manufacturing of transmission and radio relay equipment), Meissener Nachrichtentechnik GmbH (cable manufacturer); and DBP Telekom, Dresden (fibre in the local loop—OPAL project).

The visits were of varying standards. One visit was conducted in German, which proved a little disappointing for the large number of non-German speaking members of the party. The OPAL project visit was extremely well organised and informative, and consisted of a series of lectures and site visits.

Social and Cultural Events

Apart from the technical and personal networking aspects of the congress, culture has always been a high priority for FITCE. The choice of cultural venues such as the Semper Opera, Kulturpalast, Zwinger, and Dresden's many museums was a credit to the skill of the organising committee.

FITCE General Assembly

The outgoing President of FITCE, Prof. Dimitri Kouremenos, handed over the Presidency to Dr. José Luis Rojo Serrano. In addition to the reports by the General Secretary and the Treasurer, the general assembly heard that AIT had resigned from FITCE. It was agreed that AIT, the French telecommunications engineering association, would be replaced by FITCE France.

The general assembly agreed that the 34th European Telecommunications Congress would take place in Bologna, Italy, and would have a theme based on network management and mobile telephony.

Conclusions

The 33rd FITCE congress was a resounding success, in its organisation, and through the technical presentations. It provided all those who attend with an excellent opportunity to not only develop their understanding of European telecommunications issues, but to build on personal networks that will be continually challenged through increased competition.

Congress Proceedings

Copies of the congress proceedings are available from the IBTE Office for loan to IBTE Members. Please telephone (0171) 356 8008.

Biography



David Corrie
Head of the BT
Telecommunications
Master's Programme

David Corrie is Head of the Telecommunications Master's Programme (formerly the Telecommunications Staff College). He is presently working on the next generation programme, based on BT's multimedia products, with the aim of developing BT's leaders around the globe. After graduating in 1979, David took up his appointment with the System X launch team. Other key assignments include information technology management, undergraduate development programmes, a major joint venture training programme, and international training support. David is a member of the Institution of Electrical Engineers, FITCE, and a Fellow of the Institute of Personnel and Development.

Journal Awards for Volume 12

Introduction

British Telecommunications Engineering is an important record by which the membership of the Institution of British Telecommunications Engineers (IBTE) and others can keep abreast of various items of interest in telecommunications.

To encourage readers in furthering the role of the *Journal*, and to give authors due recognition for outstanding contributions, the Board of Editors operates an annual award scheme. Prizes are awarded to the authors of articles which, in the opinion of the Board, demonstrate excellence in content and presentation and which enhance the quality and range of contributions published.

Each year, a prize is awarded for the best article published in a complete volume, together with a number of prizes for runners-up. This year, for the first time, the Board, recognising the large number of high-quality articles published, introduced a highly-commended category of award.

Adrian Mahon (left), winner of the Journal's Best Article award, receives his prize from Dr. Alan Rudge



Next year the Board will also be awarding prizes to authors of units in the *Structured Information Programme*.

At its fifth annual congress and dinner held on 3-4 October 1994, the prizes for Volume 12 (April 1993-January 1994) were awarded by Dr. Alan Rudge, IBTE President.

Top Award for 'Non-Engineering' Article

The prize for the best article from Volume 12 went to Adrian Mahon for his article 'Marketing Customer Service: the Differentiator for the 1990s' in the April 1993 edition. He received a crystal bowl inscribed with the IBTE's insignia, and a cash prize.

The fact that the award this year has been given to a predominantly non-engineering article is of great significance for the IBTE and its Members. It reflects the rapidly-changing commercial environment in the telecommunications industry and the Board's response to the challenges by encouraging more

broadly-based articles of wider interest to engineers.

Despite its commercial focus, Adrian's article contains some key messages for the engineering profession, because it highlights the fact that, no matter how good the product being sold, or the technologies behind it, ultimately it is the total service package perceived by customers that differentiates one supplier from another. This is increasingly true when the product features, pricing structure and underlying technologies offered by different suppliers are tending to merge.

The article provides a very readable introduction to the subject of customer service marketing and demonstrates via both theoretical models and real customer feedback how customer service can be measured and improved by concentrating on the issues of most significance from the customers' perspective. It highlights the importance of addressing all the component parts of the service delivery process and discusses how the gaps between customers' expectations and perceptions can be bridged. In summary, it illustrates the key challenges for telecommunications engineers in the 1990s.

Runners-Up

The Board of Editors also awarded two runner-up prizes for Volume 12. Each article received a cash prize, and each author a crystal goblet.

Caller Display and Call Return

The first runner-up prize was awarded to William Dangerfield, Simon Garrett and Melv Bond for their article 'Caller Display and Call Return', which appeared in the October 1993 issue.

Caller Display and Call Return are two of the network services BT is introducing into the modern digital



Simon Garrett receives a runner-up award. His co-authors—William Dangerfield and Melv Bond—were unable to attend



Runners-up (L-R) Peter Cochrane, Rob Taylor-Hendry and (far right) Kim Fisher

network. Caller Display allows the customer to see the number of the person who is calling them and Call Return allows the customer to find the number of the last caller so that they can make a return call.

In their article, the authors describe these new services, explaining the benefits to customers and the present and future changes the services will bring. The article also explains the technical aspects, such as the signalling and network components, in a clear easy-to-follow format.

Trials of the services have raised key issues about privacy and possible abuse of the information. The authors discuss customer reactions and show the steps that BT is taking to resolve the issues to meet specific needs. They also highlight how the network services would help to reduce the problem of malicious calls by showing the caller who is calling them.

The Office You Wish You Had

The second runner-up prize was awarded to Peter Cochrane, Kim Fisher and Rob Taylor-Hendry for their article 'The Office You Wish You Had', which was published in the July 1993 edition.

Life evolved over 600 million years; it has taken 100 000 years for humans to interface to their natural environment, and only 60 years for technology to create the complex, hostile environment of the modern office. Technology has outpaced biology!

This very readable article offers a view of our workplace in the next 5–10 years, examining new freedoms offered by developments such as voice control, intelligent desks, high-definition TV, optical wireless etc. to produce intuitive, user-friendly interfaces enabling humans to communicate more freely with each other. The restrictions of the computer keyboard are gone forever!

Highly Commended

Six articles were highly-commended; the authors were presented with certificates:

- Ted Rook and Steve Green for 'BT's Internal Network' (April 1993);
- John Crackett for 'ISDN in the Dealing Room' (July 1993);
- Mick Head, Paul White and Neil Morgan for 'SuperJANET—A Strategic Partnership' (October 1993);
- Dino D'Sa and Alan Lowe for 'Planning the International Network' (October 1993);
- David Greenop, Ian Pearson and Trevor Johnson for 'Broadband—Liberating the Customer' (January 1994);
- Graham Mills and Dave Griffiths for 'PC-Based Visual Communications Services' (January 1994).

Inter-Switch Audio-Cable Rationalisation

Lead-sheathed audio cables have formed the inter-switch network (termed the MU/CJ network) since the 1920s and have been used for switched telephony and private circuits (PCs). Technology changed in the late-1960s and new cables were provided with polyethylene sheaths. The cable network has been overlaid with digital technology in the form of transverse-screen and optical-fibre cables. This new overlay network and switch modernisation have resulted in all switched telephony circuits being removed from the audio-cable network. The remaining PCs are thinly distributed over the audio cables, and the number of circuits is reducing, but the cost of maintaining this lead-sheathed network is increasing.

In Northern Home Counties (NHC), an exercise to identify the old lead-sheathed cables that have a high incidence of failure and are therefore costly to maintain, has been completed. The inter-switch cable network consists of approximately 6000 cables and, of these, 750 were identified as having a poor maintenance history. These 750 cables have been barred from being used for the provision of new circuits and, with natural churn, the number of circuits on these cables is reducing. However, a rationalisation process is still needed to remove cables from the network, since cables have to be maintained until the last circuit has been removed.

NHC offload process

The offload process adopted in NHC is an eight-stage process involving cooperation between operational divisions. This process is as follows:

Feasibility The outer-core network is managed by the tactical planners who are responsible for the strategic planning of the MU/CJ audio network. When an audio cable has been identified by a specific driver for offload (see Table 1), the planners have to decide if it is feasible to remove that cable from the network. The planners determine the number of circuits involved and the availability of alternative line plant.

Prioritise When it has been agreed that offload is feasible, the cables are listed in priority order of business benefit resulting from offload. Cables that give poor customer service are given top priority.

Verification Prioritised cables are selected for circuit verification. Records databases are compared to ascertain the validity of the circuits.

Planning Alternative routings for telephony circuits, PCs, digital sections, programme circuits, local spurs, etc. are investigated and additional line plant planned and programmed if required.

Design The on-line plant allocation (OPAL) system, that is part of integrated network systems (INS), is used to design new routings and to maintain the circuit characteristics for each of the PCs requiring rearrangements.

Execution Circuit rearrangements are undertaken with the cooperation

of the customers regarding down time of the circuits involved.

Retirement When all the circuits that were verified have been offloaded, all the cable termination (MDF bar pair) fuses are removed. This is to confirm that the cable is cleared of all circuits before the cable is retired.

Recover To prevent damage to other cables, the recovery of MU/CJ cables is undertaken only if the duct space is required for the provision of new cables. Detail planners are informed that the cable is ready for recovery; once recovery has been completed the records for the cable are removed from INS.

NHC programme

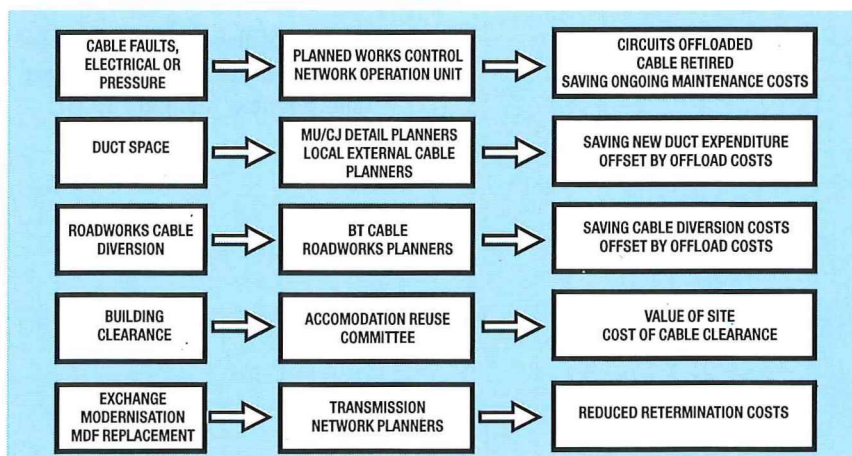
In NHC a programme to clear cables has been running since January 1993. Some MU/CJ audio cables consist of multiple sections and cables with large numbers of sections are divided into parts for offload. Some 124 cables/parts of cables have been cleared in NHC, resulting in considerable savings of expenditure from air-cylinder activities, cable length renewals, building closures, new duct provision and roadworks cable diversion.

Digital private circuit network

The KiloStream network has been enhanced to cater for the FDM closure project and the provision of all new PCs over 25 km in length. Analogue PCs provided on the KiloStream network are delivered analogue over digital (A over D) by using analogue network terminating units (ANTUs) at each end of the circuit. However, for circuits under 25 km, the first choice of routing is still audio copper pairs or PCM. Further enhancement of the KiloStream network is required to enable circuits under 25 km to be provided A over D. It is unlikely that this enhancement will provide enough capacity for the complete offload of the MU/CJ audio network, and thus it is more likely that a partial offload of cables with specific drivers will be implemented as described earlier.

Peter Legg for IBTE Thames Valley Centre

Figure 1—Drivers for initiating clearance of MU/CJ audio cables



Advanced Telecommunications Services Demonstrated in the Hague

BT and a number of European partners have developed a prototype telecommunications network of the future. Project TRIBUNE is a test-bed that contains both an exchange and a wide variety of sophisticated customer terminals, which together are capable of testing and demonstrating many of the advanced services in which telecommunications providers are taking an increasing interest.

TRIBUNE makes use of the broadband integrated services digital network (B-ISDN), which is capable of providing an almost unlimited number of telecommunications services to public and private customers. Familiar services, such as telephony and facsimile, could coexist with the more demanding ones, such as multimedia and TV broadcast services—all over a single line connection between the customer and the network.

One of the principal objectives of the project is to develop and trial real equipment and systems, designed according to the latest international B-ISDN standards and to demonstrate how interworking with other sophisticated networks can be achieved in practice. Connections are in place from the test-bed to standard office local area networks (LANs), to the existing public narrowband ISDN and to other asynchronous transport mode (ATM) pilot test-beds throughout Europe.

As a demonstrator, TRIBUNE has a planned programme of public demonstrations designed to exhibit the services of tomorrow on the equipment of today.

As a test-bed, TRIBUNE is able to offer its services to third-party equipment suppliers who may wish to connect their terminal equipment to a real service provider, offering real services over a real network—all developed in close alignment with the emerging international standards.

As a development platform, TRIBUNE is a unique study vehicle for ATM technology—a technique that is rapidly becoming accepted worldwide as the common signalling,

switching and transmission system for networks of the future.

This collaborative project between BT and a number of telecommunications equipment suppliers, R&D institutes and analysts throughout Europe is partly funded by the European Union.

First UK/France Videoconference Test via ATM Networks

In August 1994, BT and France Telecom held the first videoconference test between the UK and France using the European asynchronous transfer mode (ATM) Pilot Trial network.

This initial connection was the start of a series of trial connections to test the interworking and performance aspects of the ATM network. The connection was the culmination of a team effort between BT and France Telecom engineers, and demonstrated Europe's lead in supporting multimedia services over ATM networks.

The connection was made, using international standard H.320 videoconference terminals, between the BT Laboratories at Martlesham Heath and CNTE Laboratories at Lannion. The connection route was Martlesham-London-Paris-Lannion, using ATM cross-connect switches and a combination of 34 Mbit/s and 155 Mbit/s ATM links. The videoconference terminals' native bitstreams of 2 Mbit/s were adapted to ATM cells by using ATM Adaptation Layer Type 1.

BT Gets Smart for New Generation Phonocards

BT has placed multi-million pound contracts for new phonocards and phonocard payphones, based on smart card integrated microchip technology, to be introduced into service during 1995.

Three companies, GPT, Landis & Gyr and Schlumberger, will provide payphones able to accept smart phonocards; two companies, Gemplus and GPT, will produce the phonocards.

The new smart phonocard will be a prepaid card similar to the existing optically-encoded card and will be used in a similar manner. The cost of

the call will be debited from the value stored on the integral chip in the card.

The company believes that this latest evolution of its payphones will ensure the continued reliability of its payphone service. It is also the first step to offering advanced services for customers in the future, including the potential for using the card for payment for other services.

David Bell, BT Payphones' New Business Division Manager, said: 'BT's existing payphone technology has served the company well, but its future development is limited. Phonocards are, however, extremely popular with customers, Phonocard retailers and Phonocard collectors.

'New technology will not only improve BT's service to customers, but offer tremendous potential to make exciting new services available from BT payphones in the future'.

The new smart-technology cards will be the same size and shape as the existing BT Phonocard, the main difference being in appearance, with the gold-coloured microchip positioned on the left of the card.

The payphones will be conventional in appearance and offer a cash and card option or a card-only option.

BT Opens up ISDN for Small and Medium Businesses

BT has lowered the entry level on ISDN 30 from 15 to six channels to allow small and medium businesses anywhere in the UK to take advantage of advanced digital communications.

The new six-channel entry also meets the requirements for peripheral sites of larger companies.

BT states that small businesses can now receive the same high-quality communications as larger organisations for which ISDN 30 is already preferred for PBX connection and, increasingly, data communications.

The six-channel entry level is ideal for companies upgrading or changing their key system or PBX switch.

- Companies can now run applications such as multimedia, videoconferencing and data transfer through the PBX as a result of the emerging international standard

interface I.420 being available on an increasing number of switches. Previously this would have required direct ISDN 2 lines.

- ISDN 30 and a PABX can also provide business features such as direct dialling in, calling line identification, voice messaging and call centre service with automatic call distribution telemarketing packages.
- ISDN 30 also gives businesses extra security with options such as diversion on failure or diverse routing.

The six-channel entry will appeal to companies requiring high-quality videoconferencing through channel aggregation. Instead of investment in multiple BT ISDN 2 connections and terminal adapters, customers can use an ISDN access controller and six or more ISDN 30 channels.

All-Optical Network First for BT

BT has demonstrated the full capability of an all-optical network delivering the high-quality high-speed communications services that residential and business customers can expect in the next century.

By giving customers direct instant access to the vast fibre spectrum, an all-optical network will offer more than 1000 times the information capacity of today's integrated services digital network (ISDN), and provide the most revolutionary change in network technology since the Strowger switch was invented. It will also enable all customers to utilise the high-quality multimedia services of the future.

The multiplexed network for distributive and interactive services (MUNDI) allows customers direct access to the narrowband and broadband wavelengths of the optical spectrum in a shared fibre passive optical network (PON), and moves part of the network switching function from BT to the customer. This simplifies operations and reduces the amount of interface hardware.

Switching is achieved in two integrated ways by using tuneable wavelength transmitters and receivers at customers' terminals. The central exchange allocates wavelength channels (between 30 and 40 fully equipped) to customer terminals for the duration of each call. These are then separated out by high-density wavelength demultiplexers and route optically through an exchange switch to the desired downstream PON.

This form of distributed wavelength/space/wavelength switching leads to a simpler network, which also requires less switching hardware overall.

Unlike the broadband cable-TV networks being installed currently, the MUNDI solution provides a far more responsive and flexible network with return channels of equal bandwidth. This not only allows switched two-way broadband services, such as video telephony, videoconferencing and teleworking instantly to the home or the office, but also offers the potential for any customer to become an information provider, offering enormous growth potential for new business opportunities.

Syntegra Wins £200 000 British Gas Order

Syntegra, the systems integration business of BT, has won a £200 000 order from British Gas TransCo to develop a system which will improve efficiency by electronically linking it with its customers.

TransCo, one of the five new business units within the newly restructured British Gas, handles the transportation and storage of gas for the Public Gas Supply and Contract Trading businesses of British Gas, as well as for independent shippers.

TransCo, which generates a vast amount of meter, billing and other information for its customers, considered it was particularly critical to move from a paper-based system to an electronic system in response to the increasing liberalisation of the UK gas market. At present, one paper invoice can be several feet high, and the volume of paper can only increase with deregulation as more suppliers enter the market.

The new electronic system will improve efficiency by giving direct faster links to the customers. TransCo's customers will also be able to make fuller use of the information than ever before because it can be fed directly into their own systems where it can be further integrated and analysed.

The new system links TransCo's DEC VAX-based billing server over ISDN with shippers' own IT systems. Syntegra will provide the ISDN links and life-cycle support, as well as develop software to allow integration with TransCo's DEC system. TransCo will supply its customers with a PC which will connect into the network.

Although the system will initially provide a one-way flow of traffic between TransCo and its customers, TransCo plans to develop the system to allow a two-way flow of traffic.

Tele Danmark and Telecom Finland Join Norwegian Telecom as BT Global Partners

Tele Danmark and Telecom Finland have selected BT as their global partner in the Nordic region.

A Memorandum of Understanding (MoU) between the parties has been signed to enable Tele Danmark and Telecom Finland to offer their customers an enhanced range of communications solutions within the Nordic area and globally.

The agreement reinforces the April announcement of the strategic alliance with Norwegian Telecom by enabling the respective parties each to offer a wider and consistent domestic coverage in Denmark and Finland as well as across the Nordic area.

At the same time, Norwegian Telecom, BT and Tele Danmark announced their intention to work towards providing a wide range of telecommunications services in Sweden, building upon their existing operations in that country.

The main aim of the four parties in respect of the Nordic area is to meet the growing needs of their business customers for consistent and integrated pan-Nordic and global solutions. The agreement covered by the MoU will enable the provision of a

common portfolio and capability via their existing operations.

Initial service offerings to be made available to the region will include voice and data communications services from BT's Concert portfolio of global network services. The alliance will further aim to develop new services and/or include other services of the parties to meet specific identified needs of the Nordic business market.

DealerStream Private Circuit Services Available Throughout Central London

BT's DealerStream range of private circuit packages offering enhanced support and maintenance back-up is now available across the entire inner London 071 area.

Businesses throughout central London can choose from BT's analogue and digital circuits, backed by the DealerCare service plan. DealerStream was developed for city-based organisations and provides high-quality transmission of voice, data and image, plus increased levels of support and maintenance.

Companies no longer need 30 circuits to opt into a DealerStream package as the level of eligibility has been lowered to just 30 channels. This means, for example, a company can qualify with a single 2 Mbit/s MegaStream digital circuit equivalent to 30 channels, or a combination of analogue and KiloStream digital circuits carrying 30 channels or more.

As part of the agreement, circuits are installed within ten working days and restoration of service is guaranteed in less than five working hours of a fault being reported, an important factor for companies which rely on fast response times.

As well as DealerCare, which provides maintenance back-up between 7.00 am and 7.00 pm, Monday to Friday, DealerStream has a range of additional service options.

Businesses with fibre end-to-end circuits can opt for the increased protection of DealerCare Plus which guarantees restoration of service in less than two hours, and Premium Provision which offers delivery of circuits by the end of the next working day.

NHS to End Paper Chase with Syntegra's New Nationwide Electronic Message Service

After a year of competition, Syntegra, the system integration business of BT, has been awarded a framework contract to install and manage a new nationwide electronic Message Handling Service (MHS) for the National Health Service.

The NHS handles more than a billion paper messages each year and says that the sheer volume is stifling the potential gains from internal reform. Part of a seven-year investment programme, the new service may eventually link thousands of GP and dental practices and hundreds of hospitals and Health Authorities across the country.

In addition to creating substantial cost savings, it will speed up the transfer of patient and other information and will permit the instant distribution of emergency medical messages. Other potential users include the Dental Practices Board.

With the new service, the NHS will have one of the lowest tariffs for messaging of any major organisation in Europe.

The contract has been awarded to Syntegra, leading a team that comprises some 25 suppliers, including Digital, IBM, ICL, MDIS, Oracle, Reuters, Siemens Nixdorf and Sun Microsystems, as well as many suppliers of X.400 communications systems and specialist medical applications, such as VAMP and PROMPT. The breadth of suppliers in the Syntegra team, which between them currently provide over 50% of the IT system in the NHS, will ensure that NHS users will have a wide choice of hardware and software options.

The contract is another example of the outsourcing trend, with Syntegra responsible for the operational and financial risk of providing the message handling service for the NHS. For Syntegra, it follows other similar contracts, including, most recently, one worth more than £41M to develop and manage the CHIEF import and export system and network for HM Customs and Excise.

The NHS contract requires Syntegra to provide the basic messaging and electronic data interchange (EDI) service, and hardware and software to enable NHS sites to link into the service, and the full range of systems integration, consultancy and training services to assist NHS users in linking their existing applications into the backbone.

The service could eventually handle more than one million messages a day and will use the most up-to-date messaging technology, Digital's Mailbus400, which is soon to be used for BT's own public messaging service.

Pilot projects will be developed in Birmingham, Exeter, Ipswich and Preston before the full range of services is made available to NHS users across England from early 1995.

Protection Against Loss of Business Now Available With CircuitSure

Companies can protect key operational sites from the risk of severing their communications links in the event of fire, floods or lightning strikes with the launch of CircuitSure—BT's nationwide protection plan for digital private circuit services.

Unforeseen disasters can affect a company's daily operations and result in considerable loss of revenue, but subscribers to CircuitSure can be assured in the knowledge that their key sites are protected. CircuitSure will switch all the private circuits from the protected site to a pre-defined alternative location, which can be another of the company's sites or that of a third party, within 24 hours. When the company's original site is operational again, full service can be restored within the same time period. A premium option to switch the circuits in under four hours is also available.

The range of BT's digital private circuits covered by CircuitSure is KiloStream, KiloStream N and MegaStream which are used to support, for example, local area networks, bulk file transfer and videoconferencing. International digital circuits can also be switched

from one UK site to another, giving companies the added security of being able to maintain their overseas links.

Although BT's digital network is inherently resilient, a company should have a back-up plan to re-route their digital private circuits if a key site were hit by a disaster.

New Phone Codes Switched On

August 1994 marked the most significant milestone in the biggest revamp of the UK's telephone numbering system for more than 25 years.

From 1 August, the new national codes began to work alongside the old to allow businesses to start phasing in changes to their systems in preparation for Phoneday on 16 April 1995.

The National Code Change was announced by OFTEL two years ago in response to a growing telecommunications industry need for more codes and numbers. The communications explosion has created a vast array of new services, and liberalisation of the UK telecommunications market has seen a host of new operators enter the UK, all generating a huge demand for codes and numbers.

The extended parallel running of the new and old codes gives companies almost nine months to make changes to, for example, telephones, fax machines, switches, call loggers, call barrers, modem, security systems and databases and to phase in stationery, signage and livery changes.

BT has contacted all customers with BT supplied and maintained equipment to advise them on what needs to be done. Where engineering work is required, BT is carrying out the work free of charge. All call barring requiring engineering work was completed by BT well ahead of the 1 August date, when it became possible to make national and international calls using the new codes. As part of the European programme to rationalise access, the international code for making a call from the UK is being changed from '010' to '00'.

BT Launches New Cordless Switch, Companion 100A

Nearly three-quarters of all office calls never reach the person for whom they are intended. To alleviate both the frustration and the potential loss of business caused by lost or unwanted calls, BT is launching the Meridian Companion 100A, adding to its range of cordless switches.

The BT Companion 100A system uses small fully-featured cordless handsets and base stations to let staff make and take calls while they are away from their desks, reducing the number of unanswered calls. Base stations are installed around the work premises creating interlocking cells of radio coverage. These give the handset users complete mobility, while making and receiving calls, without loss in line connection quality.

The system is ideal for environments such as offices, factories, shops, hospitals and schools. For example, warehouse staff need the mobility to go about their work but still remain in contact. The Companion 100A will process a customer's call regardless of where a recipient is within the building.

Suitable for medium to large businesses, this new cordless switch supports up to 80 handsets and site coverage with up to 71 base stations in 24 radio cells. Using Common Air Interface and CT2 standards, the second generation in digital cordless technology, Companion 100A is linked to the company's switch allowing the handsets to act as telephone extensions. Cordless handsets can then access the same features as hard-wired telephones, including call forwarding, call transfer and conferencing. Companies do not have to replace their PBX, as the Companion 100A can be connected to any switch with analogue extensions.

The quality and combination of handsets and base stations will depend on the size of the company's premises and the number of handset users, so the price of a system can vary. The Companion 100A opens up the cordless market to larger businesses who want to enhance their customer service levels and call handling.

Syntegra and LBID Lead New Partnership Trend

A new kind of strategic relationship with customers was heralded by Syntegra, the systems integration business of BT, when it formed a partnership with Lloyds Bank Insurance Direct (LBID).

In an increasingly competitive marketplace, organisations need to react more rapidly, and the trend is towards outsourcing with the right technology partner that can adapt to meet the changing requirements of business without tying clients inextricably to a single product range or rigid method of working.

In banking, shopping and particularly insurance, the public are showing a clear and increasing preference to do business direct by telephone. Already 20% of household and motor insurance policies are sold this way, and this proportion is expected to double over the next five years. Achieving this kind of growth is impossible without highly effective and adaptable systems for handling customers' calls. Thus, for LBID, choosing the right customer contact system was of major importance as the Syntegra system is the centrepiece of LBID's contents insurance operation.

book review

Number, Please! is a new book by BT's Dave Occomore. It traces the history of the telephone system in London from the days of Thomas Edison and Alexander Graham Bell to the take over of the National Telephone Company by the General Post Office.

The book's 96 pages include many previously unpublished photographs of the period from the BT Museum and BT Archive. Chapters include descriptions of what it was like to work on overhead cables and wires that ran across the rooftops of London and life as an operator in the often-cramped switchrooms of the city.

The book, priced £9.95, will be available in late November from good booksellers or direct from the publisher, Ian Henry Publications Ltd., 20 Park Drive, Romford, Essex RM1 4LH (please add £0.61 for postage). The edition is to be limited to 600 copies.

Cerplex Acquires BT's Repair Business

The Cerplex Group Inc, one of the top high-technology logistics and repair services organisations in the US, has acquired BT's repair business comprising the assets of BT Repair Services and several other smaller repair centres. The new company, a wholly-owned subsidiary of The Cerplex Group Inc, has been named Cerplex Ltd. The principal operation will be based in Enfield, North London, with a specialised repair facility in Leeds.

Cerplex will employ some 475 people engaged on the repair and calibration of more than 650 000 items each year for BT covering a wide range of network, payphone and customer equipment. Under the agreement with BT, Cerplex Ltd. will continue to provide BT equipment repair services.

Multimedia Service for London

Videotron Corporation, London's largest cable operator, has announced plans to offer communications services in the City of London and the City of Westminster to both residential and business customers.

Under the terms of the licence awarded by the DTI earlier this year, Videotron can offer telephony services to businesses and residents. Unlike BT, it can also offer television services to businesses.

Videotron's new business TV and information service will also include video, voice, text, data and graphics to business customers all via a single connection to their desktop PC.

In the first phase, available early 1995, business customers will have access to a brand-new Reuters business news service created for Videotron, plus cable and satellite business channels, CNN, NBC SuperChannel, EuroNews and the Parliamentary Channel in addition to BBC1, BBC2, ITV and Channel 4.

These services will be complemented by real-time information on shares, commodities, currencies,

futures and news wires appearing in Windows format on the PC screen allowing users to access and work on other programmes simultaneously.

In a second phase of development, anticipated later in 1995, Videotron will extend the service to include further business TV channels, a news clipping service, travel information and reservation functions for aircraft, trains and hotels, as well as PC videoconferencing.

BA Takes a Competitive Swipe

Skyphone and British Airways have signed a multi-million pound deal to offer long-haul passengers a full suite of onboard interactive communications services. The agreement covers some 80 aircraft and is expected to deliver annual call revenues to Skyphone in excess of £17M.

Skyphone telephony will give BA passengers credit card swipe direct dialling to more than 200 countries worldwide, regardless of the aircraft's location. Skyphone consortium member BT has developed the software platform on which BA's planned on-line entertainment services will run. The consortium will be responsible for the end-to-end coordination of the BA service which will include shopping, gaming, video games, broadcast news and destination information.

Full passenger trials start early in 1995 on a single aircraft flying on key routes to South Africa, the Far East and the east coast of the USA. Subject to final approval by BA's board of directors, the Skyphone service will then be rolled out across the airline's long-haul fleet from mid-1995 onwards.

Pricing and Interconnection

In most European countries, telecommunications pricing is inflexible, unbalanced, obsolete and unsustainable. In the next five years, putting this right will be a key objective and a major challenge for all operators, but will create major difficulties for regulators, according

to The Yankee Group Europe's report 'Tariffing in the Age of Abundance'.

The report identifies eight key technological changes underpinning this revolution: digital transmission, switching and compression, asynchronous transfer mode (ATM) switching, wireless communications, computerised operations, administration and management (OA&M), the spread of new applications and services, and network computing.

Along with these, there has been significant regulation reform. The European Community's directives have been met with some success, in particular by mandating cost-oriented and non-discriminatory provision of leased circuits.

The report cautions that the going is about to get tougher for regulators: many of the commercial imperatives now driving tariff innovation—packaging, discounting, special offers, demand-based pricing and negotiated one-off deals—run counter to CEC legislation. Controlling the activities of major players will become increasingly difficult as tariffing gets more complicated.

Equally, establishing fair rules for interconnecting competing networks, and policing those agreements, will test regulators to the limits.

Discounts and packaging are clearly in the interests of some users, but present major dilemmas to regulators. In particular, predatory and discriminatory pricing designed to freeze out competition are major dangers, and is likely to be a dominant issue for the next decade.

Liberalising Satellite Communications Markets

The European Commission has given a boost to the European satellite market by adopting a Directive to liberalise satellite telecommunications equipment and services throughout the EU. The Directive should reduce the costs of deploying and operating satellite networks, facilitate the establishment of pan-European satellite networks and simplify licensing and registration. Amending a 1990 Directive on competition in telecommunications

services, the new measure concerns in particular very small aperture satellite terminals (VSATs) as well as the larger satellite dishes used for news gathering and mobile communications. Owing to liberalisation, a ten-fold increase in the volume of satellite communications is expected by the year 2000.

Unity Plan Announced for the Engineering Profession

Sir John Fairclough, FEng, Chairman of the Engineering Council, has announced a proposed plan to unite Britain's engineering profession.

The proposals have been developed by a policy group that was set up in September 1994 after the publication of the report 'Engineering into the Millennium'. The proposals will be considered by the two bodies concerned—The Engineering Council and the Council of Presidents.

The policy group was united in the belief that the proposals would forge a partnership between all those who sought to represent and promote the profession of engineering. The proposed plan provides the opportunity for leadership and vision for the profession as a whole.

Under the plan, a new body would be formed to bring together the views of the institutions, industry and academe and would provide a single voice for the profession. The new body would consist of a democratically-elected senate with 24 members elected by institution councils from their past and present council members, 24 members elected by the engineers and technicians registered with The Engineering Council and six members appointed by the Privy Council.

The senate would be the focal point of the engineering profession. It would discharge its responsibilities through two boards: one for the regulation of the profession, the other for the promotion of the profession. Members of the two boards would be selected from the senate.

Sir John said: 'The new arrangements will provide the leadership needed to take engineering into the next century and I commend them to all those who are concerned about

society and the role that engineering has to play in it.'

Sir John said that it was hoped to have the new arrangements in place by mid-1995.

NVQs are Britain's 'Best-Kept Secret'

The education and training of engineers is outdated and not strategically focused to suit industry's needs, The Engineering Council's Assembly at Churchill College, Cambridge, has been told.

Speaking at the assembly in July, Graham Mackenzie, Director General of the Engineering Employers' Federation, said that an overhaul of the process to ensure that engineers and technicians are multi-skilled, have business and finance skills and are good communicators is essential.

He told 200 delegates and observers representing the UK's 290 000 professional engineers and technicians that industry wanted a continuous and flexible structure which 'does not draw lines where none can be discerned particularly in relation to the qualities required of engineers and technicians.'

He said that continuing professional development (CPD) was an integral part of an engineer's career progression.

Paul Shepherd, Chairman and Managing Director, Shepherd Construction, said that the introduction of National Vocational Qualifications (NVQs) has encouraged the education and training process at further education colleges to become more attuned to the demands of business.

'NVQs provide potentially sound progress from existing qualifications but they are not going to solve our selection problems,' he said.

'For NVQs to achieve their full potential they will need to be assessed at the workplace. Yet increasingly I feel that this process is being overtaken by the colleges, often through the use of simulation exercises. This is not a criticism of the colleges but more of a challenge for industry,' he said.

John Berkeley, Manager, Education and Careers, Rover Group, said: 'As long as NVQs remain one of Britain's few surviving well-kept secrets, GNVQs are regarded with less esteem than A-levels and degree courses fail to make explicit what graduates can actually do, the incoherence of the present structure will continue to inhibit performance.'

'We need a new vision of education—business collaboration, a coherent and integrated post-14 credit framework capable of supporting seamless continuity of lifelong learning and a re-engineered process for developing the full potential of the entire workforce within a changing world of employment arrangements and relationships.'

Dr David Fussey, Vice Chancellor at the University of Greenwich and a member of The Engineering Council, told the assembly that the idea of inventing a 'Super CEng' grade for registration had been well aired but the consensus was in favour of retaining the existing grades of Chartered Engineer (CEng), Incorporated Engineer (IEng) and Engineering Technician (EngTech), but raising educational standards for each.

Dr. Fussey said that for the CEng grade, a qualification beyond the BEng is needed, which implied either continuing professional development or an MEng, each equivalent to a further year of full-time study. Both the IEng and EngTech grades would be similarly supplemented by further full-time study.

Call for Better Environmental Design

Existing standards of design of goods for the telecommunications industry fall short of current ideas of sustainability. This is the principal conclusion of a report from the RSA Environmental Design Workshops—'Ecodesign in the Telecommunications Industry'.

The intention of the report is to stimulate further debate on environmental design issues within the telecommunications industry and beyond. This is the first report of its

kind in the UK and offers a basic primer on the subject both for telecommunications specialists and for those seeking to extend the thinking to other trade sectors.

The main outcomes of the workshop consisted of urgent proposals for practical projects. These included the establishment of a European environmental network group within the telecommunications industry and the development of an environmentally-friendly telephone for home use.

Some proposals are already being implemented. For example, the European Institute for Research and Strategic Studies into Telecommunications (EURESCOM) is supporting a project 'Telecommunications and the Environment' which will consider the telecommunication industry's contribution to sustainable development. This has support from BT, Telia AB and Telefónica. The report outlines these projects and includes papers by several experts in the field, recommendations of working groups, a bibliography and a list of participants in the

workshop. (RSA: 8 John Adam Street, London WC2N 6EZ. Telephone: (0171) 930 5115.)

ATM Fears Threaten Superhighway Growth

Plans to build ATM-based European information superhighways are at risk because conservative telecommunications operators (TOs) are afraid of the impact that ATM will have on their revenues. They are holding back despite the enthusiasm of equipment manufacturers and potential new network operators.

Research by Cambridge-based telecommunications strategy consultants Analysys Ltd., in a new report 'ATM: Strategies, Benefits and Technology', shows that many of Europe's TOs are being very slow in moving to deploy ATM in their networks. Susan Ablett, Analysys Senior Consultant, said: 'They are afraid that ATM could shrink their existing revenue base because of its ability to integrate voice, data and image into a single network and

deliver existing services to customers more cheaply than at present. But ATM also brings opportunities, in the shape of advanced services, which will generate new sources of revenue for the TOs.'

Analysys has developed two scenarios to portray the development of the ATM superhighway, and the information, communications and entertainment (ICE) market it will support, over the years to the end of the century.

'All the players will base their strategies for the next five years or so on which scenario they think is more likely to develop,' says Susan Ablett. 'The report explores the implications of whether they predict the future correctly or not.'

New President for IEE

The new President of the Institution of Electrical Engineers is Professor Sir David Davies, Chief Advisor, Ministry of Defence. He replaces Dr. Alan Rudge, BT Board Member for Technology.

XI International Symposium on Subscriber Loops and Services

CALL FOR PAPERS

The overall theme of ISSLS '96, which is being held in Melbourne, Australia, from 4-9 February 1996, is

Building Towards a Broadband World

Exciting narrowband and broadband access capabilities, supported by the accelerating convergence of CATV, wired and wireless communications, broadcasting and interactive multimedia, are leading to challenging business opportunities. ISSLS '96 will highlight building access networks and systems which make good business sense, and will consider the interplay between access technology possibilities and social change.

The ISSLS symposium, with its excellent reputation and tradition for quality and stimulating technical papers and discussions, has become the focal point for setting directions for access networks and services evolution. ISSLS '96 in Melbourne will provide a platform for customers, service providers, access network operators and equipment suppliers to engage in fruitful dialogue.

Papers are invited discussing the following broad topic areas:

- Customer access—business opportunities and risks.
- Access operations and management.
- Social change, competitive forces or technology—what is driving change in access.
- Network and technology—leaps in access, CATV, CPE and associated systems.

For further information about ISSLS '96, the symposium theme and the procedure for submitting papers, please write to the following address, call or fax: ISSLS '96 Conference Secretariat, IREE Australia, PO Box 79, Edgecliff NSW 2027, Australia. Tel: +61-2-327-4822. Fax: +61-2-362-3229.

BRITISH TELECOMMUNICATIONS ENGINEERING ISSN 0262-401X

Published in April, July, October and January by *British Telecommunications Engineering Journal*, 2-12 Gresham Street, London, EC2V 7AG. (Formerly *The Post Office Electrical Engineers' Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any statements made nor the opinions expressed in any of the articles or correspondence in this *Journal*, unless any such statement is made specifically by the Board.

© 1994: The Institution of British Telecommunications Engineers.

Printed in Great Britain by Unwin Brothers Limited, The Gresham Press, Old Woking, Surrey, GU22 9LH.

Subscriptions and Back Numbers

Price: Single copy—companies, universities, libraries and other bodies £9.25 (£10 including postage UK, £11.25 including postage overseas); private individuals (paid by personal cheque) £5.00 (£5.75 including postage UK, £7.00 including postage overseas). Annual subscriptions (including postage)—companies etc. £40.00 UK, £45.00 overseas; private individuals (paid by personal cheque) £23.00 UK, £28.00 overseas.

Overseas customers can pay by sterling drafts drawn on London banks or by sterling Eurocheques.

Price to BT staff: £2.25 per copy.

Back numbers can be supplied if available at the single copy rates given above.

Orders, by post only, should be addressed to *British Telecommunications Engineering Journal* (Sales), PP: G012, 2-12 Gresham Street, London EC2V 7AG.

Remittances for all items (except binding) should be made payable to 'BTE Journal' and should be crossed ' & Co'.

Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, *British Telecommunications Engineering Journal*, PP: G012, 2-12 Gresham Street, London EC2V 7AG. (Telephone: (0171) 356 8050. Fax: (0171) 356 7942.)

Communications

All communications should be addressed to the Editorial Office, *British Telecommunications Engineering Journal*, PP: G012, 2-12 Gresham Street, London EC2V 7AG. (Telephone: (0171) 356 8050. Fax: (0171) 356 7942.)

Copyright

The entire contents of this *Journal* and the *Structured Information Programme* are covered by general copyright and special permission is necessary for reprinting long extracts, but editors are welcome to use not more than one-third of any article, provided that credit is given at the beginning or end, thus: 'From British Telecommunications Engineering'.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by the *British Telecommunications Engineering Journal* for users registered with the Copyright Clearance Centre's (CCC's) Transactional Reporting Service, provided that the base fee of \$5.00 per copy is paid directly to CCC, 27 Congress Street, Salem, MA 01970, USA. For those organisations that have been granted a photocopy licence by CCC, a separate system of payment has been arranged. Payment is additionally required for copying of articles published prior to 1978.

0262-401X/94 \$5.00 + .00



THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

IBTE Local-Centre Contacts

The following is a list of Local-Centre contacts to whom enquiries about the Institution may be directed. Where no Membership Secretary details are shown, membership enquiries should be directed to the Local Secretary.

Centre	Local Secretary	Membership Secretary
East Anglia	Terry Birdseye (01702) 373330	Keith Phillips (01206) 717244
East Midlands	Dave Bostrom (01908) 656215	Ian Bethell (0115) 947 8587
Lancs & Cumbria	Trevor Hughes (01254) 666261	
Liverpool	John Johnson (01925) 415705	Clive Maude (0151) 229 4074
London	Jim Skinner (0171) 728 7015	Peter Martin (0171) 492 6138
Manchester	David White (0161) 200 9231	Mike Ku (0161) 200 9343
Martlesham Heath	Peter Shaw (01473) 644378	Ian Gilbert (01473) 645500
North Downs	Nick Smith (01227) 474823	
North East	Richard Jackson (0191) 261 3560	
North Wales & Marches	Peter Clay (01743) 274205	
Northern Ireland	Allan MacDougall (01232) 215301	Brenda Sharvin (01232) 320092
Scotland East	Graham Neilson (0131) 668 5878	Lewis Shand (0131) 668 5593
Severnside	Phil Reynolds (0117) 920 5224	Dean Shephard (0117) 920 6457
South Downs	Graham Lindsay (01273) 222633	
South Midlands	Tony Atkinson (01234) 277874	Jim Coley (01234) 274840
South Wales	Peter Coleman (01222) 514524	Keith Voyce (01222) 555075
Staffordshire	Ray Piper (01785) 813483	Paul Field (01785) 762419
Thames Valley	Stephen Graham (01252) 337104	Dave Prout (01734) 501244
West Midlands	Richard Harwood (01527) 550252	Paul Jervis (0121) 230 4181
West Scotland	Graham Wellby (0141) 220 3997	Jill Maxwell (0141) 220 5837
Westward	Rob Rand (01392) 267114	Chris Gould (01392) 267113
Yorks & Lincs	Steve King (0113) 237 8124	Paul Horncastle (0113) 276 0242

Associate Section Zone Committee Contacts

The following is a list of Associate Section Zone contact points to whom enquiries about the Associate Section should be directed.

Zone Committee	Contact	Telephone No.
London	Terry McCullough	(0181) 456 6798
Midland	John Sansom	(0116) 251 3155
North East	Geoff Jenkinson	(0113) 237 8163
Northern Home Counties	Andy Edmonds	(01473) 645171
Northern Ireland	Brian Walker	(01232) 232413
North West	Phil Holt	(0161) 600 8111
Scotland	Bill Lindsay	(01382) 302629
Southern Home Counties	Martin Harvey	(01424) 754385
South West	John Horton	(0117) 932 9303
Wales	Ferg Brown	(01952) 249886

IBTE National HelpLine — (0171) 356 8008

FEDERATION OF TELECOMMUNICATIONS ENGINEERS
OF THE EUROPEAN COMMUNITY (FITCE)

34th European Telecommunications Congress

CALL FOR PAPERS

The management of networks and services has always been a vitally important topic for telcos. However, recently the information technology evolution has allowed interconnection and interworking among different databases as well as between network systems, thus causing dramatic changes in the management of telecommunications. The importance of management systems greatly increases in an international competitive scenario, in which the competitive margin relies much more on the productive process rather than the product itself.

In this context, the 34th FITCE Congress, which will be held in Bologna, Italy, from 10–16 September 1995, will be based on the theme:

'Telecommunications Management'

It will focus on the management aspects (OA&M) of present and future networks and services, with particular reference to the system aspects of:

- integrated management of different network elements having different functions and supplied by different manufacturers;
- integrated management of different services;
- quality supervision and control in the operators' and customers' view.

FITCE is inviting challenging papers on this theme, covering one or more of the following aspects:

- economic and marketing aspects, including billing and pricing;
- relationships among equipment providers, network providers, service providers and end users, considering present trends and the satisfaction of user needs;
- standardisation, normalisation and specification:
 - marketing approach in a competitive environment, and
 - national and international aspects in a multi-vendor environment;
- software engineering implications, and experiences in other fields;
- present initiatives at the national, European and global levels.

If you are interested in submitting a presentation, please prepare a summary of no more than 250 words giving a clear indication of the theme and coverage of the proposed paper. The summary, which should be prepared on the standard FITCE form, should include details of your full name, place and date of birth, job function or title, company, business address and business telephone and facsimile numbers. The summary should be sent before **24 February 1995** to:

Paul Nichols, FITCE UK Papers Coordinator, Post Point G012, 2–12 Gresham Street, London EC2V 7AG (Telephone: (0171) 356 8022; Facsimile: (0171) 356 7942).

Copies of the FITCE standard form are available from the above address, and will be sent on request. Authors will be advised of the papers that have been selected during April 1995, after the final programme has been determined. The full text of a selected paper, which is required by 26 May 1995, should be about 2000 words to allow for a 25 minute presentation and 5 minutes for questions and debate.



Contents

Guest Editorial: The Need for an Architectural Approach	185
Building on the Architectural Framework	186
A New Structure for London's Public Switched Telephone Network	192
User-Friendly Weapons for the Competitive Fight	201
New Network Infrastructures for the 21st Century Part 2—Ultra-fast self-routing networks	206
The Future of Learning	213
Distributed Management of Future Global Multi-Service Networks	221
Planning the Concert Packet Services Network	227
Ensuring Quality in Design of Customer Networks	233
The Design and Implementation of the BOXER Network	240
The 'Value' of the Telecommunications Engineer in Eastern and Western Europe	246
33rd European Telecommunications Congress	250
Journal Awards for Volume 12	254

